BioLuminate 1.3

User Manual

Schrödinger Press

November 2013

Contents

Document Conventions ix

Chapter 1: Introduction 1

1.1 **Overview of BioLuminate Features** 1
1.1.1 Protein Analysis Tools 1
1.1.2 Protein Structure Tools 2
1.1.3 Peptide Tools 2
1.1.4 Protein Alignment and Docking Tools 3
1.1.5 Protein Mutation and Modification Tools 3
1.1.6 Antibody Tools 4
1.2 **Running Schrödinger Software** 4
1.3 **Starting Jobs from the BioLuminate Interface** 6
1.4 **Citing BioLuminate in Publications** 7

Chapter 2: The BioLuminate Interface 9

2.1 **The Main Window** 9
2.2 **Structures and Projects** 11
2.3 **The Toggle Table** 12
2.3.1 Quick Access Buttons 12
2.3.2 Table Rows 13
2.3.3 The Action Menu 14
2.3.3.1 Changing the View 14
2.3.3.2 Minimizing the Structure Energy 15
2.3.3.3 Changing the Appearance of Structures 15
2.3.3.4 Displaying Polar Contacts 16
2.3.3.5 Generating an Atom Selection 17
2.3.3.6 Displaying Atoms Related By Crystallographic Symmetry 18
2.3.3.7 Modifying an Atom Selection 18
2.3.3.8 Removing Selections 20
2.3.3.9 Renaming Rows 20
2.3.3.10 Duplicating Rows 20
2.3.3.11 Removing Entries from the Workspace 20

2.3.3.12 Deleting Entries from the Project 20
2.3.3.13 Creating Project Entries 21
2.3.3.14 Adding and Removing Hydrogens 21
2.3.3.15 Removing Waters 21
2.3.3.16 Computing Properties 21
2.3.4 The Show Menu 22
2.3.5 The Hide Menu 23
2.3.6 The Label Menu 24
2.3.7 The Color Menu 25
2.4 Shortcut Menus 26

Chapter 3: Analyzing Protein Quality **27**

3.1 Ramachandran Plot 27
3.2 Protein Report 28
3.3 Plot Toolbar 30

Chapter 4: Analyzing Residue Properties **31**

Chapter 5: Identifying Consensus Molecules **35**

5.1 Choosing the Target Protein 36
5.2 Finding and Aligning Homologs 36
5.3 Viewing Consensus Molecules 37

Chapter 6: Identifying Reactive Residues **39**

Chapter 7: Predicting Aggregation Regions **43**

7.1 Creating an Aggregation Surface 43
7.2 Setting Options for Aggregation Surfaces 45
7.3 Analyzing the Surface 45
7.4 Using the Results for Mutation Studies 47

Chapter 8: Analyzing Protein Interface Interactions 49

8.1 Selecting the Interacting Proteins 50

8.2 Running the Analysis 50

8.3 Examining the Results 52

8.4 Viewing Interaction Properties in the Workspace 53

Chapter 9: Locating Large-Scale Motions 55

Chapter 10: Predicting Peptide Properties 57

10.1 Alpha Helix Stability 57

10.2 QSAR from Sequences 59

10.2.1 Loading the Sequences and Properties 60

10.2.2 Building a QSAR Model 61

10.2.3 Examining the Model 62

10.2.4 Applying the Model to Other Sequences 64

10.2.5 Setting Options for PLS Methods 64

10.2.5.1 PLS Options 64

10.2.5.2 KPLS Options 65

10.3 Peptide Binding to a Receptor 67

10.3.1 Defining the Protein Receptor Region 67

10.3.2 Specifying the Peptides to Dock 68

10.3.3 Setting Docking Options 69

10.3.4 Running the Job 70

Chapter 11: Protein-Protein Docking 71

11.1 Preparation of Proteins for Docking 72

11.2 Docking a Protein to Another Protein 72

11.3 Applying Constraints 74

11.4 Creating Dimers and Trimers 75

11.5 Docking an Antigen to an Antibody 75

Chapter 12: Homology Modeling of Proteins 77

Chapter 13: Residue and Loop Mutation 81

13.1 Residue Mutations 82
13.1.1 Selecting the Residue to Mutate 83
13.1.2 Defining the Mutation 83
13.1.3 Refining the Mutated Structure 84

13.2 Insertions, Deletions, and Loop Swaps 85
13.2.1 Choosing the Original and Replacement Loop 87
13.2.2 Editing the Replacement Loop 88
13.2.3 Choosing the Output and Method 89

Chapter 14: Cross-Linking Proteins 91

14.1 Preparing the Proteins for Cross-Linking 91

14.2 Picking the Connection Residues 92

14.3 Defining the Linkers 94

14.4 Choosing Methods and Running the Job 95

14.5 Examining the Results 95

Chapter 15: Scanning for Residue Mutations 97

15.1 Selecting and Analyzing the Protein 97

15.2 Setting Up and Running the Job 98
15.2.1 Choosing Residues to Mutate 98
15.2.2 Choosing the Mutations 99
15.2.3 Setting Optimization Options for Affinity Maturation 101
15.2.4 Choosing Refinement Options 102
15.2.5 Running the Job 103

15.3 Using Homologs for Identifying Mutations 104
15.3.1 Obtaining Homologs from a BLAST Search 104
15.3.2 Importing Homologs 104
15.3.3 Aligning Homologs 105

15.3.4 Selecting Residues by Homology ... 106
15.3.5 Selecting Residues by Structural Attributes ... 107
15.3.6 Making the Selection ... 107

15.4 Examining the Mutation Results ... 107
15.4.1 Binding Affinity Prediction ... 110
15.4.2 Stability Prediction ... 110

Chapter 16: Locating Possible Mutations for Disulfide Bridges ... 111

16.1 Selecting and Analyzing the Protein ... 111

16.2 Choosing Residue Pairs to Mutate ... 113

16.3 Relaxing the Structure Around the Mutation Site ... 115

16.4 Running the Job ... 115

16.5 Examining the Results ... 116

16.6 Workflow Summary ... 118

Chapter 17: Antibody Modeling ... 119

17.1 Modeling an Antibody Structure ... 119
17.1.1 Importing the Antibody Sequence ... 120
17.1.2 Choosing a Database ... 121
17.1.3 Selecting the Coordinates for the Framework Region ... 122
17.1.4 Filtering the Database by Property ... 123
17.1.5 Generating the Loop Model ... 125
17.1.6 Selecting Clusters for a Loop ... 126
17.1.7 Refining Loops ... 128
17.1.8 Summary ... 130

17.2 Humanizing Antibodies by Residue Mutation ... 131
17.2.1 Analyzing the Antibody ... 131
17.2.2 Finding Homologs ... 132
17.2.3 Selecting Regions to Mutate ... 133
17.2.4 Setting Up Residue Selection Criteria ... 133
17.2.5 Selecting the Residues and Their Mutations ... 134

17.2.6 Setting Refinement Options for Mutated Residues 135
17.2.7 Running the Mutation Job 136
17.2.8 Analyzing the Results 136
17.2.9 Summary 138
17.3 Humanizing Antibodies by CDR Grafting 139
17.4 Antibody Databases 142
References **145**
Getting Help **147**

Document Conventions

In addition to the use of italics for names of documents, the font conventions that are used in this document are summarized in the table below.

Font	Example	Use
Sans serif	Project Table	Names of GUI features, such as panels, menus, menu items, buttons, and labels
Monospace	`$SCHRODINGER/maestro`	File names, directory names, commands, environment variables, command input and output
Italic	*filename*	Text that the user must replace with a value
Sans serif uppercase	CTRL+H	Keyboard keys

Links to other locations in the current document or to other PDF documents are colored like this: Document Conventions.

In descriptions of command syntax, the following UNIX conventions are used: braces { } enclose a choice of required items, square brackets [ ] enclose optional items, and the bar symbol | separates items in a list from which one item must be chosen. Lines of command syntax that wrap should be interpreted as a single command.

File name, path, and environment variable syntax is generally given with the UNIX conventions. To obtain the Windows conventions, replace the forward slash / with the backslash \ in path or directory names, and replace the $ at the beginning of an environment variable with a % at each end. For example, `$SCHRODINGER/maestro` becomes `%SCHRODINGER%\maestro`.

Keyboard references are given in the Windows convention by default, with Mac equivalents in parentheses, for example CTRL+H (⌘H). Where Mac equivalents are not given, COMMAND should be read in place of CTRL. The convention CTRL-H is not used.

In this document, to *type* text means to type the required text in the specified location, and to *enter* text means to type the required text, then press the ENTER key.

References to literature sources are given in square brackets, like this: [10].

Introduction

1.1 Overview of BioLuminate Features

BioLuminate offers a wide range of tools for protein modeling, protein engineering, protein analysis, peptide analysis, and antibody modeling. In addition to the unique tools, BioLuminate provides access to many of the related tools in the Schrödinger software suite. The BioLuminate interface is a customization (profile) of the Maestro interface, designed for protein modeling.

This manual documents the unique tools and capabilities of BioLuminate, and provides references to other documents for the related tools. A brief description of the tool set is given below, with links to the relevant parts of this manual or of other manuals. The descriptions are classified by function. These tools are divided between the Tools menu, where the action does not take much time and may be interactive, and the Tasks menu, where a job may need to be run that takes a larger amount of time.

1.1.1 Protein Analysis Tools

The protein analysis tools provide information on a protein and its properties. No change is made to the protein structure.

- Protein Structure Quality Viewer (Tools → Protein Structure Quality): Show reports on deviations of protein parameters from standard values, in graphical and tabular form. See Chapter 3.
- Residue Analysis (Tasks → Residue Analysis): Calculate energetic and other properties of residues. See Chapter 4.
- Consensus Visualization (Tools → Protein Consensus Viewer): Locate consensus waters, counter ions and ligands in a set of homologs to a reference protein. See Chapter 5.
- Reactive Protein Residues (Tools → Reactive Residue Identification): Identify residues that are prone to specified reactions, by matching sequence patterns and some structural information. See Chapter 6.
- Aggregation Surface (Tasks → Aggregation Surface): Predict regions on a protein surface that have a propensity for aggregation. See Chapter 7.
- Protein Interaction Analysis (Tasks → Protein Interaction Analysis): Analyze the interactions at the interface of two proteins. See Chapter 8.

- Low Mode Vibrational Sampling (Tasks → Low Normal Mode Analysis): Locate and visualize large-scale vibrational motions in a protein. See Chapter 9.
- SiteMap (Tools → Binding Site Identification): Locate druggable sites on a protein. See the *SiteMap User Manual*.

1.1.2 Protein Structure Tools

The protein structure tools allow you to fix structures from the PDB or other sources that are missing information needed for modeling or are missing atoms, predict the structure of proteins by homology modeling, and predict the structure and stability of alpha helices in small peptides.

- Protein Preparation Wizard (Tools → Protein Preparation): Prepare proteins for modeling by assigning bonds, fixing structural defects, removing unwanted parts, assigning protonation and tautomeric states, and refining the structure. See the *Protein Preparation Guide*.
- Simple Homology Modeling (Tasks → Simple Homology Modeling): Predict the structure of proteins using homology modeling, where the homology is high and the alignment of the query and the template is straightforward. See Chapter 12.
- Structure Prediction (Tasks → Advanced Homology Modeling): Predict the structure of single-chain or multi-chain proteins, including multimers, by homology modeling. See Chapter 3 through Chapter 5 of the *Prime User Manual*.
- Refinement (Tasks → Loop + Sidechain Prediction or)Tasks → Implict Solvent Refinement + Analysis): Refine protein structures by performing predictions of selected side chains or loops, or minimizations of various parts of protein structures. See Chapter 6 of the *Prime User Manual*.

1.1.3 Peptide Tools

The peptide tools allow you to predict various properties of small peptides from their sequences—see Chapter 10.

- Peptide Helicity (Tasks → Peptide Alpha Helicity): Predict the stability of alpha helices for small peptide sequences, using molecular dynamics.
- Peptide Docking (Tasks → Peptide Docking): Dock peptides to a receptor, starting from the sequence. The receptor is largely rigid, the conformational space of the peptide is explored.
- Peptide QSAR (Tasks → Peptide QSAR): Predict properties of small peptides using a QSPR (sequence-property) model based on peptide descriptors.

1.1.4 Protein Alignment and Docking Tools

The alignment tools include tools that structurally superimpose two proteins (or structures), and a tool for docking one protein to another. These tools perform rigid-body translation of the structures to obtain the best alignment.

- Align Binding Sites (Tools → Binding Site Alignment): Align the sites on a set of proteins at which drug-like molecules can bind. See Chapter 7 of the *Prime User Manual*.
- Protein Structure Alignment (Tools → Protein Structure Alignment): Structurally align two or more proteins, using secondary structure information as well as coordinates. See Chapter 7 of the *Prime User Manual*.
- Superposition (Tools → Superposition): Align two or more structures by minimizing the RMSD of a selected set of atoms. See Section 10.3 of the *Maestro User Manual*.
- Protein-Protein Docking (Tasks → Protein-Protein Docking): Predict how two proteins interact, using a rigid body search algorithm. See Chapter 11.

1.1.5 Protein Mutation and Modification Tools

- Residue Scanning (Tasks → Residue Scanning): Systematically perform single mutations of protein residues to determine how energetic and other properties change, and to identify mutations that can effect desired changes. See Chapter 15.
- Affinity Maturation (Tasks → Affinity Maturation): Perform multiple sequential mutations of protein residues to optimize the binding affinity of two proteins or the stability of a protein. See Chapter 15.
- Cysteine Mutation (Tasks → Cysteine Mutation): Locate residue pairs that can reasonably form disulfide bonds if one or both of the residues are mutated to cysteine and perform the mutation, or locate disulfide bonds and mutate one of the residues to another type to break the disulfide bond. See Chapter 16.
- Residue and Loop Mutation (Tasks → Residue and Loop Mutation): Mutate a single residue to a standard or custom residue, invert the chirality of a residue, delete or insert multiple residues in a single loop, or swap a loop for another loop. See Chapter 13.
- Crosslink Proteins (Tasks → Crosslink Proteins): Join two protein chains from their termini with a series of linkers, to produce a set of structures. See Chapter 14.

1.1.6 Antibody Tools

BioLuminate provides a specialized set of tools for modeling antibodies, including managing databases of antibody structures, homology modeling of antibodies, antibody humanization, and antigen-antibody docking.

- Antibody Prediction (Tasks → Antibody Modeling → Prediction): Predict the structure of the CDR region of an antibody by homology modeling, using homology and database methods. See Section 17.1 on page 119.
- Antibody Humanization: Residue Mutation (Tasks → Antibody Modeling → Humanization → Residue Mutation): Humanize an antibody by performing mutation of residues selected manually or on the basis of homology to human antibodies. See Section 17.2 on page 131.
- Antibody Humanization: CDR Grafting (Tasks → Antibody Modeling → Humanization → CDR Grafting): Humanize an antibody by grafting the CDR loops on to a human antibody. See Section 17.3 on page 139.
- Antibody Database Management (Tasks → Antibody Modeling → Database Management): Select, create, and update antibody databases. See Section 17.4 on page 142.
- Protein-Protein Docking (Tasks → Protein-Protein Docking): Dock an antigen to an antibody. See Section 11.5 on page 75.

1.2 Running Schrödinger Software

Schrödinger applications can be run from a graphical interface or from the command line. The software writes input and output files to a directory (folder) which is termed the *working directory*. If you run applications from the command line, the directory from which you run the application is the working directory for the job. The BioLuminate interface is a customization of the Maestro interface. You can also use the standard Maestro as your working interface.

Linux:

To run any Schrödinger program on a Linux platform, or start a Schrödinger job on a remote host from a Linux platform, you must first set the `SCHRODINGER` environment variable to the installation directory for your Schrödinger software. To set this variable, enter the following command at a shell prompt:

csh/tcsh: `setenv SCHRODINGER` *installation-directory*

bash/ksh: `export SCHRODINGER=`*installation-directory*

Once you have set the SCHRODINGER environment variable, you can run programs and utilities with the following commands:

```
$SCHRODINGER/program &
$SCHRODINGER/utilities/utility &
```

You can start the BioLuminate interface with the following command:

```
$SCHRODINGER/maestro -profile BioLuminate &
```

It is usually a good idea to change to the desired working directory before starting the BioLuminate interface. This directory then becomes the working directory.

Windows:

The primary way of running Schrödinger applications on a Windows platform is from a graphical interface. To start the BioLuminate interface, double-click on the BioLuminate icon, on a Maestro project, or on a structure file; or choose Start → All Programs → Schrodinger-2013-3 → BioLuminate. You do not need to make any settings before starting BioLuminate or running programs. The default working directory is the Schrodinger folder in your documents folder (Documents on Windows 7/Vista, My Documents on XP).

If you want to run applications from the command line, you can do so in one of the shells that are provided with the installation and have the Schrödinger environment set up:

- Schrödinger Command Prompt—DOS shell.
- Schrödinger Power Shell—Windows Power Shell (if available).

You can open these shells from Start → All Programs → Schrodinger-2013-3. You do not need to include the path to a program or utility when you type the command to run it. If you want access to Unix-style utilities (such as awk, grep, and sed), preface the commands with sh, or type sh in either of these shells to start a Unix-style shell.

Mac:

The primary way of running Schrödinger software on a Mac is from a graphical interface. To start the BioLuminate interface, click its icon on the dock. If there is no BioLuminate icon on the dock, you can put one there by dragging it from the SchrodingerSuite2013-3 folder in your Applications folder. This folder contains icons for all the available interfaces. The default working directory is the Schrodinger folder in your Documents folder (`$HOME/Documents/Schrodinger`).

Running software from the command line is similar to Linux—open a terminal window and run the program. You can also start BioLuminate from the command line in the same way as on Linux. The default working directory is then the directory from which you start BioLuminate.

You do not need to set the SCHRODINGER environment variable, as this is set in your default environment on installation. To set other variables, on OS X 10.6 and 10.7 use the command

`defaults write ~/.MacOSX/environment` *variable* `"`*value*`"`

and on OS X 10.8 use the command

`launchctl setenv` *variable* `"`*value*`"`

1.3 Starting Jobs from the BioLuminate Interface

To run a job from the BioLuminate interface, you open a panel from one of the menus (e.g. Tasks), make settings, and then submit the job to a host or a queueing system for execution. The panel settings are described in the help topics and in the user manuals. When you have finished making settings, you can use the Job toolbar to start the job.

You can start a job immediately by clicking Run. The job is run on the currently selected host with the current job settings and the job name in the Job name text box. If you want to change the job name, you can edit it in the text box before starting the job. Details of the job settings are reported in the status bar, which is below the Job toolbar.

If you want to change the job settings, such as the host on which to run the job and the number of processors to use, click the Settings button. (You can also click and hold, and choose Job Settings from the menu that is displayed.)

You can then make the settings in the Job Settings dialog box, and choose to just save the settings by clicking Save, or save the settings and start the job by clicking Save and Run. These settings apply only to jobs that are started from the current panel.

If you want to save the input files for the job but not run it, click the Settings button and choose Write. A dialog box opens in which you can provide the job name, which is used to name the files. The files are written to the current working directory.

The Settings button also allows you to change the panel settings. You can choose Read, to read settings from an input file for the job and apply them to the panel, or you can choose Reset Panel to reset all the panel settings to their default values.

You can also set preferences for all jobs and how the interface interacts with the job at various stages. This is done in the Preferences panel, which you can open at the Jobs section by choosing Preferences from the Settings button menu.

Note: The items present on the Settings menu can vary with the application. The descriptions above cover all of the items. Jaguar has an Edit item and extra functions for the Read and Write items, which are described later in the manual.

The icon on the Job Status button shows the status of jobs for the application that belong to the current project. It starts spinning when the first job is successfully launched, and stops spinning when the last job finishes. It changes to an exclamation point if a job is not launched successfully.

Clicking the button shows a small job status window that lists the job name and status for all active jobs submitted for the application from the current project, and a summary message at the bottom. The rows are colored according to the status: yellow for submitted, green for launched, running, or finished, red for incorporated, died, or killed. You can double-click on a row to open the Monitor panel and monitor the job, or click the Monitor button to open the Monitor panel and close the job status window. The job status is updated while the window is open. If a job finishes while the window is open, the job remains displayed but with the new status. Click anywhere outside the window to close it.

Jobs are run under the Job Control facility, which manages the details of starting the job, transferring files, checking on status, and so on. For more information about this facility and how it operates, as well as details of the Job Settings dialog box, see the *Job Control Guide*.

1.4 Citing BioLuminate in Publications

The use of this product should be acknowledged in publications as:

BioLuminate, version 1.3, Schrödinger, LLC, New York, NY, 2013.

The BioLuminate Interface

The BioLuminate interface is a customized form of the Maestro interface that is specially designed for biologics use. It inherits most of the capabilities of the Maestro interface (though organized differently), and it has features of its own.

This chapter focuses on the features that are unique to BioLuminate. Summaries of the main Maestro features are given, with references to the *Maestro User Manual* for details. If you have never used Maestro, you should be able to gain a basic understanding of its operation from this chapter.

If you prefer to use the standard Maestro interface, you can do so. Most of the capabilities of BioLuminate are available from the BioLuminate submenu of the Applications menu, and some of them are on the Tools menu.

You can open the BioLuminate interface as follows:

- **Windows:** Double-click the BioLuminate icon on the desktop
- **Mac:** Go to Applications → SchrodingerSuite2013 and double-click the BioLuminate icon
- **Linux:** Start Maestro and choose BioLuminate in the Choose Profile dialog box, or use the command `$SCHRODINGER/maestro -profile BioLuminate`

2.1 The Main Window

The BioLuminate main window opens with the following features displayed by default:

- **Menu bar.** This is at the top of the window on Linux and Windows, and is the menu bar on the Mac.
- **Manager toolbar.** This toolbar is just below the menu bar on Linux and Windows, and at the top of the window on the Mac. Each label on this toolbar displays or hides another toolbar. By default they are all hidden, as much of their function is available in the Toggle Table. See Section 2.4 of the *Maestro User Manual* for details of the toolbars.
- **Toggle Table.** This dockable panel is displayed on the right side of the main window. You can undock it from the main window and redock it with the docking button.

  

  Many other panels are also dockable. You can change the docking behavior in the Preferences panel (Edit → Settings → Preferences, or CTRL+,)

- **Workspace.** This is the large black area that occupies the main part of the main window. It is where structures are displayed, along with any associated objects such as surfaces and text labels.
- **Status bar.** This bar is below the Workspace. At the left is a button that displays information on what jobs are running, which you can click to open the Monitor panel for detailed information on your jobs. When the pointer is not over an atom in the Workspace, the status bar gives information on the contents of the Workspace. When the pointer is over an atom, the status bar gives information on the identity of the atom. For more information, see Section 2.5 of the *Maestro User Manual*.

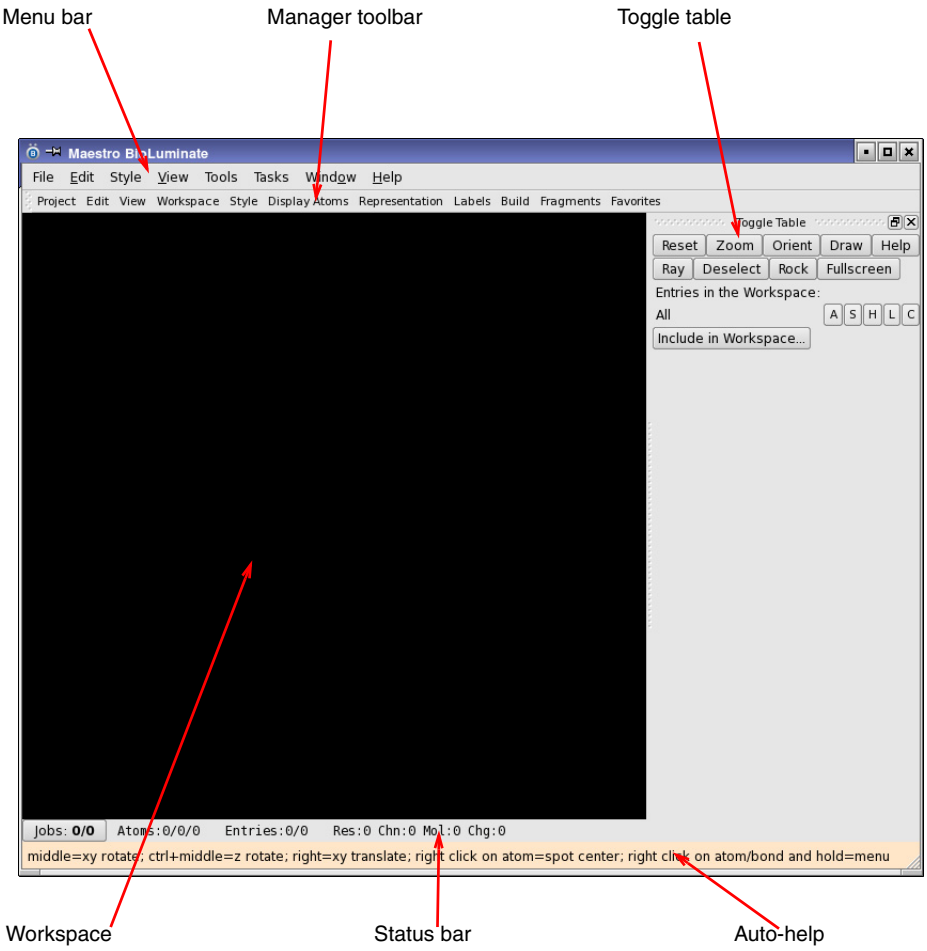

Figure 2.1. The BioLuminate main window.

- **Auto-help.** This orange-yellow bar at the bottom of the main window gives tips on the current action that can be performed in the Workspace.

There are several other components of the main window that can be displayed when needed, by choosing Edit → Settings and then choosing the component. These components include the Sequence Viewer, which displays the sequences of proteins that are in the Workspace; the Find toolbar (also opened with CTRL+F, ⌘F), which you can use to find structural components in the Workspace, like chains or residues; and the Clipping Planes window, which shows a top view of the Workspace and the planes where the structures in the Workspace are clipped for display.

2.2 Structures and Projects

When you start BioLuminate, a new, temporary project is created. Projects are the data structures that Maestro uses to store and manage molecular structures, such as proteins and ligands. Each such molecular structure in a project is stored in an *entry*. An entry can consist of multiple molecules—chains of a protein, waters, cofactors, ions, and so on. The molecular structure (coordinates, charges, and bonding information) is stored along with any properties of the structures, such as the PDB ID and crystallographic information, surfaces associated with the structure, and display information for showing the structure in the Workspace. Properties of individual atoms are stored as well as properties of the structure as a whole.

To add structures to the project, you can *import* them from an external source, such as a file or the Protein Data Bank (PDB). To import structures from a file, choose File → Import Structures. To get structures directly from the PDB (either from a local copy or from the web), choose File → Get PDB. Details of both of these methods for importing can be found in Section 3.1 of the *Maestro User Manual*.

To see a list of all the entries in the project, you can open the Project Table panel with CTRL+T (⌘T) or Window → Show → Project Table. This panel lists the entries in the project with their properties, and provides ways of doing actions on the entries and their properties, such as management tasks, sorting, grouping, plotting, import and export of entries and properties. Full details of the operation of the Project Table can be found in Chapter 9 of the *Maestro User Manual*. The menu organization in the Project Table panel in the BioLuminate interface is a little different from that in the standard Maestro interface: the Table menu in the standard interface is split between the Table menu and the Tools menu in the BioLuminate interface.

2.3 The Toggle Table

The Toggle Table panel can be used to interact with structures in the Workspace. The panel has a set of buttons at the top for quick access to some common actions. The main part of the panel consists of a row for each entry in the Workspace, with a set of buttons, or "toggles", that can be used to perform actions. These buttons are actually cascading menus, from which you can make selections. In addition to the rows for each entry, there is a row for all entries, labeled All, and rows for selected atoms in the Workspace. The entry rows are labeled with the entry title. Below the rows is a button for including entries in the Workspace.

When the toggle table is displayed, a set of shortcut (or context) menus is also available in the Workspace, which you open by right-clicking. These menus offer the same functions as the toggles.

The interaction with the Workspace provided by the Toggle Table panel is very similar to the operation of PyMOL. If you are familiar with PyMOL, this interface should be easy to learn. If you are familiar with Maestro but not with PyMOL, you can close the Toggle Table panel and use the standard Maestro interaction with the Workspace.

The features of the Toggle Table panel are described in detail in the following subsections.

Note: The terminology used in the Toggle Table panel is the PyMOL terminology, which is somewhat different from that used in standard Maestro.

2.3.1 Quick Access Buttons

This set of buttons at the top of the Toggle Table panel provides quick access to a number of actions, some of which are also on the menus.

- Reset—Reset the view to the default view, in which the view axes are aligned with the coordinate axes of the structure.
- Zoom—Change the view of the Workspace so that all atoms fit inside the Workspace area.
- Orient—Orient the Workspace structures by translating and rotating them so that the center of mass is at the origin, the largest principal axis lies on the x axis, and the second-largest principal axis lies on the y axis.

  Note: This operation changes the coordinates of the structures, not just the coordinates of the view.

- Draw—Save an image of the Workspace to a file in TIFF, JPEG, or PNG format. (Same as File → Save Image.)
- Help—Open the help topic for the panel in your browser.
- Ray—Use PyMOL to draw a ray-traced image of the Workspace. This feature requires PyMOL to be installed, and either the `PYMOL4MAESTRO` or `PYMOL_PATH` environment variable must point to the directory where the PyMOL executable file resides.
- Deselect—Clear the current Workspace selection.
- Rock—Rotate the Workspace back and forth smoothly. Click once to start rotation, click again to stop.
- Fullscreen—Switch Maestro to full screen mode. To exit full screen mode, press the ESCAPE key or click this button again.

2.3.2 Table Rows

Each row in the Toggle Table has a title and a set of five action buttons, labeled A, S, H, L, and C. These buttons open cascading menus, and are described in the following sections.

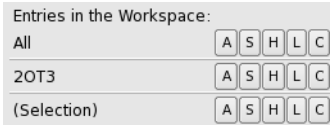

There are three distinct types of rows in the Toggle Table:

- **The All row:** Actions taken in this row apply to all entries in the Toggle Table.
- **Entry rows**: These rows apply to a single project entry that is currently in the Workspace. The name of the row is the Title property of the entry. Entry rows are deleted when the entry is excluded from the Workspace, and a new row is added when an entry is added to the Workspace.
- **Selection rows:** When a group of atoms is selected, a selection row is added to the table. Actions in this row apply to the group of atoms that was selected to create this row and even apply after those atoms are no longer selected. The selection row named (Selection) always refers to the most recent group of selected atoms. If a new selection is made while a selection row is active, that row now refers to the new set of selected atoms. Only one selection can be active at a time.

  Selection rows can be renamed: renamed selection rows always refer to the same set of atoms regardless of subsequent Workspace selections. To rename a selection row, choose A → Rename Selection. Selection row names are always enclosed in parentheses.

  Selection rows are deleted when any atoms they refer to are removed from the Workspace. To delete a selection row, choose A → Delete Selection.

Using the A button submenus, selection rows can also be duplicated, copied, and extracted to define a new entity that is independent of the objects from which the selection was originally derived.

Some operations or menu items change based on whether they are being applied to the All row, an entry row, or a selection row. The descriptions below primarily describe the behavior for entry rows. When a behavior changes for the All row or a selection row, the change is noted after the description.

Clicking the name of the All row or an entry row changes the visibility of that object in the Workspace. When the object's visibility is off, the name is dimmed. This is a quick way of showing or hiding the atoms in entries. Hiding atoms does not remove them from the Workspace, so any action taken on the entry or the entire Workspace applies to the hidden atoms as well as the visible atoms.

For instance, if the Workspace contains ten entries and only one entry is visible, choosing Clean from the A menu in the All row operates on all ten structures. This may take a considerable amount of time to complete and lead to unexpected results. Similarly, if a panel imports structures from the Workspace, it imports all ten structures rather than just the visible structure.

Clicking the name of the current selection row deselects the selected atoms, but the selection remains defined. Clicking the name of any other selection row selects the atom group that the row refers to, and any currently selected atoms that are not part of this atom group are deselected.

2.3.3 The Action Menu

The A button opens the Action menu, from which you can choose a variety of actions for the structure defined by the table row. Not all of the actions are available in every table row.

2.3.3.1 Changing the View

The first four actions on the Action menu change the view (camera angle) of the structure defined by the table row. The Workspace coordinate system has the origin in the middle of the Workspace, the *x* axis is the horizontal axis, the *y* axis is the vertical axis, and the *z* axis is coming out of the screen.

- Zoom—Change the view of the Workspace so that all the atoms in the structure fit inside and fill the Workspace area. In the All row, this action is equivalent to clicking the Zoom button at the top of the panel.
- Orient—Orient the structure by translating and rotating it so that the center of mass is at the origin, the largest principal axis lies on the x axis, and the second-largest principal axis lies on the y axis. When you apply this operation to the selection, it is the center of

mass and principal axes of the selection that are used, but the entire structure is reoriented.

Note: This operation changes the coordinates of the structures, not just the coordinates of the view (the camera angle).

- Center—Center the structure in the Workspace, by translating the structure so that its centroid is at the center of the Workspace.
- Origin—Set the center around which rotation is performed to the centroid of the structure. The centroid need not be at the Workspace origin.

2.3.3.2 Minimizing the Structure Energy

You can minimize the energy of the Workspace or the selected atoms using the OPLS_2005 force field, by choosing Clean from the Action menu. The current selection is updated (or a new selection created) to refer to the minimized atoms. To avoid starting a lengthy calculation that is better performed in the background, Clean is limited to structures of fewer than 1000 atoms.

2.3.3.3 Changing the Appearance of Structures

The set of actions on the Preset submenu change the appearance (representation) of the structures using various preset styles. Some of these representations take a little time to set up, and a progress bar is displayed at the bottom of the table.

- Simple—Show proteins as ribbons (C alpha trace) colored by chain, ligands and bound receptor as sticks, and solvent, disulfides and ions as lines. Atom colors are not changed.
- Simple (no solvent)—Same as Simple, but no waters are shown.
- Ball and Stick—Show atoms and bonds in ball and stick, with no protein ribbons.
- B-Factor—Show the protein as tubes with residues colored by the B-factors of the residues, in a relative scheme that ranges from shades of blue, through green and yellow to red.
- Technical—Show atoms as sticks, colored with rainbow colors by residue position, and show polar contacts (hydrogen bonds) as yellow dotted lines.
- Ligand—Show proteins as ribbons colored with rainbow colors by residue position, and ligands as lines. All protein atoms within 5 Å of the ligand are shown as lines, with carbons colored with rainbow colors by residue position. Waters are shown as sticks, polar contacts are shown as yellow dotted lines, and the view is zoomed in to the atoms shown.

- Ligand Sites—This submenu shows variations on the Ligand preset that alter the way the protein or region around the ligand is shown:
  - Cartoon—Show proteins as cartoons rather than ribbons.
  - Solid surface—Show the molecular surface of the protein around the ligand as an opaque surface, colored by the nearest non-hydrogen atom.
  - Transparent surface—Show the molecular surface of the protein around the ligand as a semi-transparent surface, colored by the nearest non-hydrogen atom; show atoms and bonds as sticks.
  - Dot surface—Show the molecular surface of the protein around the ligand as a dot surface, colored by the nearest non-hydrogen atom; show atoms and bonds as sticks.
  - Mesh surface—Show the molecular surface of the protein around the ligand as a mesh surface, colored by the nearest non-hydrogen atom; show atoms and bonds as sticks.
- Pretty—Show proteins as cartoons colored with rainbow colors by residue position and ligands as sticks.
- Pretty (with solvent)—Same as Pretty but waters are shown as ball and stick.
- Publication—Same as Pretty, but protein helices are two-sided.
- Publication (with solvent)—Same as Pretty (with solvent), but protein helices are two-sided.
- Protein Interface—Color ribbons and carbons by chain, show anything not at a protein interface as cartoon ribbons, and interface residues as ball and stick. Non-carbon atoms retain their previous coloring. Interface residues are residues in a chain with more than 300 atoms that are within 4.5 Å of another chain with more than 300 atoms.
- Antibody—Show everything as cartoon ribbons colored by antibody structure. The light chain is colored in red hues, the heavy chain is colored in blue hues, and everything else is colored green. Constant regions are dark hues and the CDR regions are bright hues. The light chain L1-L3 loops are shaded orange to brown, while the heavy chain H1-H3 are shaded grey-blue to cyan.
- Default—Show everything as lines with default colors (colored by element with green carbons).

2.3.3.4 Displaying Polar Contacts

You can turn on or off the display of polar contacts (hydrogen bonds) between various groups of atoms from the Polar Contacts submenu. The display is turned on with A → Polar Contacts → Find; it is turned off with A → Polar Contacts → Remove All. The atom groupings are:

- Within Object
- Involving Side Chains
- Involving Solvent
- Excluding Solvent
- Excluding Main Chain
- Excluding Intra-Main Chain
- Just Intra-Side Chain
- Just Intra-Main Chain
- To Other Atoms In Entry
- To Other Atoms In Entry Excluding Solvent
- To Any Atoms
- To Any Atoms Excluding Solvent

Each menu choice clears any previous choice before applying the new choice.

If you want more flexibility in choosing the atom groups between which polar contacts are shown, you can use the Measurements panel. Choose Style → Measurements → H-Bonds from the menu bar at the top of the Workspace to open this panel.

2.3.3.5 Generating an Atom Selection

You can select atoms in predefined groups by choosing A → Generate → Selection and then choosing the group. The Generate item is not available for the All or Selection rows. The predefined groups are:

- All—All atoms.
- Polymer—Backbone and side-chain atoms.
- Organic—Ligand atoms.
- Solvent—Water atoms.
- Surface Residues—All residues with solvent-exposed surface area greater than 10 $Å^2$.
- Protein Interface—Residues in a chain of more than 300 atoms that are within 4.5 Å of another chain of more than 300 atoms.

Atom selections can also be generated by picking atom groups in the Workspace. To set the kind of atom group you want to pick, choose Edit → Pick Mode then the atom group name (Atoms, Residues, Chains, Molecules, or Entries), using the menu bar at the top of the Workspace. You can also set the mode by typing the first letter of the name when the pointer is in the Workspace. To pick an atom group, click on an atom in the Workspace that belongs to the group. You can see information about the atom in the Status bar when you pause the pointer over the atom.

2.3.3.6 Displaying Atoms Related By Crystallographic Symmetry

If you want to see atoms from nearest-neighbor crystal symmetry mates of your structure, you can choose A → Generate → Symmetry Mates. To show symmetry-related atoms requires crystal symmetry information to be present for the structure. You should also ensure that you have only the structure that you want to see the related information for, because this action applies to the entire Workspace.

The Symmetry Mates submenu items control the cutoff distance from the original structure for which symmetry mate atoms should be displayed. Note that no matter how far the cutoff is placed from the original structure, only the nearest-neighbor mates are created and shown. The symmetry mates are created as separate, temporary entries ("scratch entries") and are shown in the toggle table. You can remove the symmetry mates by choosing Generate → Symmetry Mates → Show None.

2.3.3.7 Modifying an Atom Selection

The Modify item is available for selection rows only, and replaces Generate on the Action menu. It allows you to alter the group of selected atoms to include or exclude other atoms. After a Modify action, the selection row applies to the new group of atoms.

The Modify actions all have a choice of atom groups to which they apply.

- Around—select all atoms or residues within a given distance from the current set of atoms, and deselect the current set of atoms. The distance is chosen from the submenu, and can encompass atoms only or be filled to entire residues that have any atoms within the chosen distance.
- Expand—Expand the current selection to include all atoms or residues within a given distance from the current set of atoms. The distance is chosen from the submenu, and can encompass atoms only or be filled to entire residues that have any atoms within the chosen distance.
- Extend—Expand the current selection to include all atoms or residues within a given number of bonds from the current set of atoms.
- Invert—Deselect the current set of atoms and select all other atoms within a given atom group. The atom groups can be chosen from the submenu:
    - Within Objects—All atoms in the entry that are not part of the selection are selected.
    - Within Chains—In each chain that contains selected atoms, the unselected atoms are selected, and the selected atoms are deselected.
    - Within Residues—In each residue that contains selected atoms, the unselected atoms are selected, and the selected atoms are deselected.

  - Within Molecules—In each molecule that contains selected atoms, the unselected atoms are selected, and the selected atoms are deselected.
  - Within Any—All atoms in the entire Workspace that are not part of the selection are selected.

- Complete—Add all other atoms within a given atom group to the selection. The atom groups can be chosen from the submenu.
  - Residues—In each residue that contains selected atoms, all atoms are selected.
  - Chains—In each chain that contains selected atoms, all atoms are selected.
  - Objects—In each entry that contains selected atoms, all atoms are selected.
  - Molecules—In each molecule that contains selected atoms, all atoms are selected.
  - C-alphas—All alpha carbons for residues within the selection are selected. All other atoms are deselected.
- Restrict to—Reduce the selection to only those atoms within a specific group that are currently selected. The available atom groups are:
  - Object—Restrict the selection to atoms in a specific entry, chosen from the submenu.
  - Selection—Restrict the selection to atoms in a specific selection row, chosen from the submenu.
  - Visible—Restrict the selection to atoms that are visible.
  - Polymer—Restrict the selection to backbone and side-chain atoms.
  - Organic—Restrict the selection to ligand atoms.
  - Solvent—Restrict the selection to water atoms.
  - Inorganic—Restrict the selection to atoms other than H, C, N, O, F, P, S, Cl, Br, or I.
- Include—Include additional atom groups in the current selection. The available atom groups are:
  - Object—Include atoms in a specific entry, chosen from the submenu.
  - Selection—Include atoms in a specific selection row, chosen from the submenu.
  - Visible—Include all visible Workspace atoms.
- Exclude—Exclude specific atoms from the selection. The available atom groups are:
  - Object—Exclude atoms in a specific entry, chosen from the submenu.
  - Selection—Exclude atoms in a specific selection row, chosen from the submenu.
  - Visible—Exclude atoms that are visible.
  - Polymer—Exclude backbone and side-chain atoms.
  - Organic—Exclude ligand atoms.
  - Solvent—Exclude water atoms.
  - Inorganic—Exclude atoms with element other than H, C, N, O, F, P, S, Cl, Br, or I.

2.3.3.8 Removing Selections

For the All row, Delete Selections removes any selection rows from the Toggle Table. All atoms are deselected but are otherwise unaltered.

For selection rows Delete Selection just removes the row from the Toggle Table. The atoms in this selection group are deselected but otherwise remain unaltered.

2.3.3.9 Renaming Rows

You can change the name of an entry row, and thereby change the Title of the project entry with the Rename action. A text box is displayed instead of the name, and you can type a new name in the box. If you decide you do not want to change the name after all, press ESC.

You can change the name of a selection row with the Rename Selection action. If you do this for the default selection row, the selection is preserved for future use as a named selection. You can also rename named selection.

This command not available for the All row.

2.3.3.10 Duplicating Rows

You can duplicate table rows with the Duplicate action, with the exception of the All row. The action is different for entry rows and for selection rows.

For entry rows, the action creates a new project entry below the entry that is the duplicate of the entry. Both structures remain in the Workspace and are listed in the Toggle Table.

For selection rows, this action creates a duplicate selection row in the Toggle Table. No project entry is created. This can be useful if you want to use a selection as the basis for another selection.

2.3.3.11 Removing Entries from the Workspace

To remove (exclude) an entry from the Workspace, choose A → Remove from Workspace for that row. Removing an entry from the Workspace just means that it is no longer in your working area. The entry remains in the project, and is listed in the Project Table. You can add it back to the Workspace by using the Include in Workspace button or the Project Table.

To clear the Workspace entirely, choose Remove Everything from Workspace in the All row.

2.3.3.12 Deleting Entries from the Project

To remove an entry from the project, choose Delete from Project in the entry row. The structure is removed from the Workspace and from the project. It no longer appear in the Project Table.

2.3.3.13 Creating Project Entries

You can create new project entries either from a selection or from an entry row, by choosing Copy to New Project Entry. The new project entry is created by copying the selected atoms or entries, and it is added to the Workspace. This command is the same as Duplicate if you choose it for an entry row.

To create a project entry by removing atoms from entries and placing them into a new entry, you can choose Extract to New Project Entry in a selection row. The new project entry contains the atoms in the selection, and the atoms are deleted from their current structure.

2.3.3.14 Adding and Removing Hydrogens

You can add or remove hydrogen atoms from a row by choosing Hydrogens → Add or Hydrogens → Remove. If you perform this action in the All row or an entry row, the atoms are removed from the structure. If you perform this action in a selection row, the hydrogens in the selection are removed, or they are added to complete the valences of the atoms in the selection.

You can also add hydrogens from the main menu bar with Edit → Add Hydrogens or from the Edit toolbar.

2.3.3.15 Removing Waters

You can remove waters from structures with the Remove Waters action. If you perform this action in the All row or an entry row, the atoms are removed from the structure. If you perform this action in a selection row, the waters in the selection are removed from the structure.

2.3.3.16 Computing Properties

To compute simple properties of the object (selection or entry), you can use the Compute action. The properties you can compute are:

- Atom Count—The number of atoms.
- Formal Charge Sum—Sum of atom formal charges.
- Partial Charge Sum—Sum of atom partial charges.
- Surface Areas—Compute surface areas. Solvent-Accessible surface area (SASA) uses a solvent radius of 1.4 Å while Molecular surface area uses the same algorithm but with no solvent to expand the atomic radii. The surface area is computed for the object in the context of the visible atoms in the Workspace.

The computed properties are displayed in a window that opens, and you can copy and paste the text. They are not stored in the Project Table.

2.3.4 The Show Menu

The S button opens the Show menu, which contains commands that alter the display of structures in the Workspace. There are three different types of display for structures:

- Atomic representations such as lines, sticks, ball and sticks or spheres
- Ribbon representations such as ribbons and cartoons
- Surface representations such as solid or mesh surfaces

Each atom may have one of each representation type active at once, but cannot have multiple representations of the same type active. For instance, an atom can be shown simultaneously with lines, cartoon ribbons and a mesh surface. However, an atom cannot be shown simultaneously by both ball and stick and sphere representations because they are both atomic representations. If an action is chosen to show an atom with a representation in the same category as an existing representation for that atom, the existing representation is removed and the atom is shown with the new representation. For instance, if an atom is shown as lines, and the Ball and Stick menu item is chosen, the lines representation is removed and the atom is shown with ball and stick representation.

The common representations that can be applied are:

- Lines—Set the atomic representation to lines (wire frame).
- Sticks—Set the atomic representation to sticks (tube).
- Ball and Stick—Set the atomic representation to ball and stick.
- Ribbon—Set the ribbon representation to ribbon (CA trace tube).
- Cartoon—Set the ribbon representation to cartoons.
- Label—Show any labels defined for this structure.
- Nonbonded—Set the atomic representation of any atom with no attachments (bonds) to Lines (wire frame). These atoms appear as small stars. No change is made to atoms that have attachments (bonds).
- Spheres—Set the atomic representation to spheres (CPK). The sphere radii are the van der Waals radii of the atoms.
- Nonbonded Spheres—Set the atomic representation of any atom with no attachments (bonds) to spheres (CPK). No change is made to atoms that have attachments (bonds).

These items appear on both the Show menu itself and on the As submenu. When you choose from the Show menu, the representation is added to the display. When you choose from the As submenu, the previous representation is replaced by the new choice.

For instance, a residue shown with lines and cartoon ribbons is shown as only ball and stick if S → As → Ball and Stick is chosen, but is shown as cartoon and ball and stick if S → Ball and Stick is chosen.

Two commands on the menu for entry rows and the All row create and display molecular surfaces. The surface is created if it does not exist, otherwise the color and representation of the surface is changed.

- Mesh—Display a mesh molecular surface colored by current atom color.
- Surface—Display a solid molecular surface colored by current atom color.

The remaining four items have submenus from which you can set the representation to lines, sticks or spheres:

- Organic—Set the atomic representation of backbone and side-chain atoms.
- Main Chain—Set the atomic representation of backbone atoms.
- Side Chain—Set the atomic representation of side-chain atoms.
- Disulfides—Set the atomic representation of disulfide atoms.

2.3.5 The Hide Menu

The H button opens the Hide menu, which contains commands that hide features in the Workspace. To redisplay these features, use the Show menu. Each item hides only the particular feature for the row, and leaves any other features as they are. Hiding atoms also hides any representation of the bonds to those atoms.

- Everything—Hide all features: atomic, ribbon, and surface representations and labels.
- Atoms—Hide atoms and bonds.
- Ribbon—Hide all ribbon representations.
- Cartoon—Hide all ribbon representations.
- Label—Hide labels. The labels remain defined and can be redisplayed.
- Nonbonded Atoms—Hide atoms with no attachments, such as Cl^- ions.
- Mesh—Hide surfaces.
- Surface—Hide surfaces.
- Main Chain—Hide backbone atoms.
- Side Chain—Hide side-chain atoms.
- Waters—Hide water atoms.
- Hydrogens—Hide nonpolar or all hydrogens, as chosen from the submenu.

- Symmetry Mates—Removes all crystal symmetry mates from the Workspace. This is a Workspace setting, so affects all symmetry mates in the Workspace.
- Polar Contacts—Remove polar contact (hydrogen bond) markers.
- All Others—Hide everything for all atoms in the Workspace other than the atoms defined in the row.

2.3.6 The Label Menu

The L button opens the Label menu, which contains commands that label features in the Workspace. Unless otherwise specified, labels are created for every atom the row applies to. Atom labels created by these commands are not cumulative. Any existing atom labels are removed when new ones are created.

Labels can be cleared from the Workspace with L → Clear. They can be hidden with H → Label. If they are hidden, they can be redisplayed with S → Label, while if they are cleared, they need to be created (usually by other Label menu commands) before they can be shown again.

Residues—Label the first carbon atom in each residue with the three-letter PDB code and residue number.

Chains—Label the first and last residue in each chain with the chain name.

The next set of commands offers a choice of the label content, including identifiers and numeric properties.

- Atom Name—Label each atom with its PDB atom name.
- Element Symbol—Label each atom with its element symbol.
- Residue Name—Label each atom with the three letter PDB code of the residue it is in.
- Residue Identifier—Label each atom with the residue number of the residue it is in.
- Chain Identifier—Label each atom with the name of the chain it is part of.
- B-factor—Label each atom with its PDB B-factor value, if it exists.
- Occupancy—Label each atom with its partial occupancy data if it exists.
- VDW Radius—Label each atom with its van der Waals radius.
- Other Properties—Submenu with other properties that can be used for labels:
  - Formal Charge—Label each atom with its formal charge.
  - Partial Charge (0.00)—Label each atom with its partial charge to two decimal places.
  - Partial Charge (0.0000)—Label each atom with its partial charge to four decimal places.
  - MacroModel Text Type—Label each atom with its MacroModel atom type.

  - MacroModel Numeric Type—Label each atom with the numerical index for the MacroModel atom type.
  - Stereochemistry—Label each atom with E,Z and R,S stereochemistry.
- Atom Identifiers—Submenu of atom identifiers that can be used for labels.

2.3.7 The Color Menu

The C button opens the Color menu, which contains commands to color the atom representations by various coloring schemes including by element, chain, substructure, B-factor and entry. The schemes are grouped into classes, each of which is represented by a menu item.

- Color by Element—Color H, C, N, O and S atoms. Default colors are white for H, green for C, blue for N, red for O and yellow for S. There are several choices for modifying the color scheme on this submenu.
  - Reset HNOS—Set H, N, O and S atoms to their default color. Carbon atoms remain their current color.
  - Custom Color {C}HNOS—Pick the color for carbon atoms and set H, N, O, and S atoms to their default color. Clicking on the menu item opens a palette of colors to choose from for carbon atoms, while selecting Recent Color Choices lists the most recent colors chosen by this command.
  - Custom Color {H}CNOS—Pick the color for hydrogen atoms and set C, N, O, and S atoms to their default color. Clicking on the menu item opens a palette of colors to choose from for hydrogen atoms, while selecting Recent Color Choices lists the most recent colors chosen by this command.
- Color by Chain—Color atoms by chain:
  - by Chain (Carbons)—Change the color of carbon atoms only.
  - by Chain (Calpha)—Change the color of alpha carbon atoms only.
  - by Chain—Change the color of all atoms.
  - Chainbows—Each chain is colored with rainbow colors.
- Color by Substructure—Helices, sheets and loops are colored by the chosen color scheme. All other atoms retain their current color.
- Color by Spectrum—Color all residues by a spectrum of colors.
  - Rainbow (Carbons)—Color carbon atoms only with rainbow colors by residue position. The chain is divided into segments, each of which has residues of the same color.
  - Rainbow (Calpha)—Color alpha carbon atoms only with rainbow colors by residue position.

    - Rainbow—Color all atoms with rainbow colors by residue position.
    - B Factors—Color atoms by the residue B factor.
    - B Factors (Calpha)—Color alpha carbon atoms only by the residue B factor.
- Auto—Color groups of atoms via a cycle of colors:
    - Carbons—color carbon atoms by the next color in the cycle.
    - All Atoms—color all atoms by the next color in the cycle.
    - Carbons by Object—color carbon atoms a different color for each entry.
    - All Atoms by Object—color all atoms a different color for each entry.
- Custom Color All Atoms—Change the color of all atoms to a single custom color. Clicking on the menu item opens a palette of colors to choose from, while selecting Recent Color Choices lists the most recent colors chosen by this command.

2.4 Shortcut Menus

The Workspace has two shortcut (context) menus when the Toggle Table panel is open. One is for the selected atoms, and one is for the entire Workspace. To show the shortcut menu for the Workspace, right-click and hold in an empty part of the Workspace. To show the shortcut menu for the selection, right-click and hold over one of the selected atoms. If you right-click and hold over an unselected atom, the selection changes: it's the same as picking that atom first, then right-clicking and holding on it.

Most of the menu items are the same as those on the Toggle Table buttons or button menus. The Selection shortcut menu has a Disable item, which turns off the selection. The Workspace shortcut menu has Enable and Disable actions, which you can use to display or undisplay any of the Workspace rows. These actions are also available from some of the submenus. This menu also has a Select action, which you can use to create a selection from the visible (displayed) atoms. It also allows you to operate on the visible atoms only or on all atoms in the Workspace.

Analyzing Protein Quality

The quality of a protein structure is often measured by deviations from values reported in the PDB. You can analyze a protein and display tabular and graphical reports on its quality in the Protein Structure Quality Viewer panel, which you open by choosing Tools → Protein Structure Quality.

If there is a protein in the Workspace, it is analyzed when you open the panel. Otherwise, you can display a protein in the Workspace and click Analyze Workspace to perform the analysis.

At the top of the panel, the protein table lists the chains in the protein that is analyzed along with various measures of the overall structure quality. You can analyze multiple proteins and they are all listed in the table, and you can select multiple chains in a single protein for reporting, but you cannot select multiple proteins.

The remainder of the panel consists of two tabs that show different data: the Ramachandran Plot tab, and the Protein Report tab.

3.1 Ramachandran Plot

The Ramachandran Plot tab displays a Ramachandran plot of the dihedral angles. Glycines are plotted as triangles, prolines are plotted as squares, all other residues are plotted as circles. The plot has three regions: favored, allowed, and disallowed, which are colored by the color scheme selected from the option menu at the bottom of the tab. There are two schemes: Green/Cyan/Red and Red/Yellow/White, for coloring the favored, allowed and disallowed regions.

Pausing the cursor over a point displays information for that residue at the top of the panel, and highlights the residue in the Workspace. Clicking on a point selects the point and zooms the Workspace image in to that residue, and highlights it with pale yellow markers. The point is displayed as an outline instead of solid black. The residue information is displayed at the top of the panel. Click again on the point to deselect it.

The plot area has a toolbar, which is described in Section 3.3.

Below the plot, you can select options to change the appearance of the residues in the Workspace structure for each of the three regions. The appearance is changed by coloring the residues and modifying the molecular representation. The coloring is applied when you analyze the Workspace or change the color scheme, so you should set these options first, and then click Analyze Workspace or change the color scheme. Deselecting any of these options does not

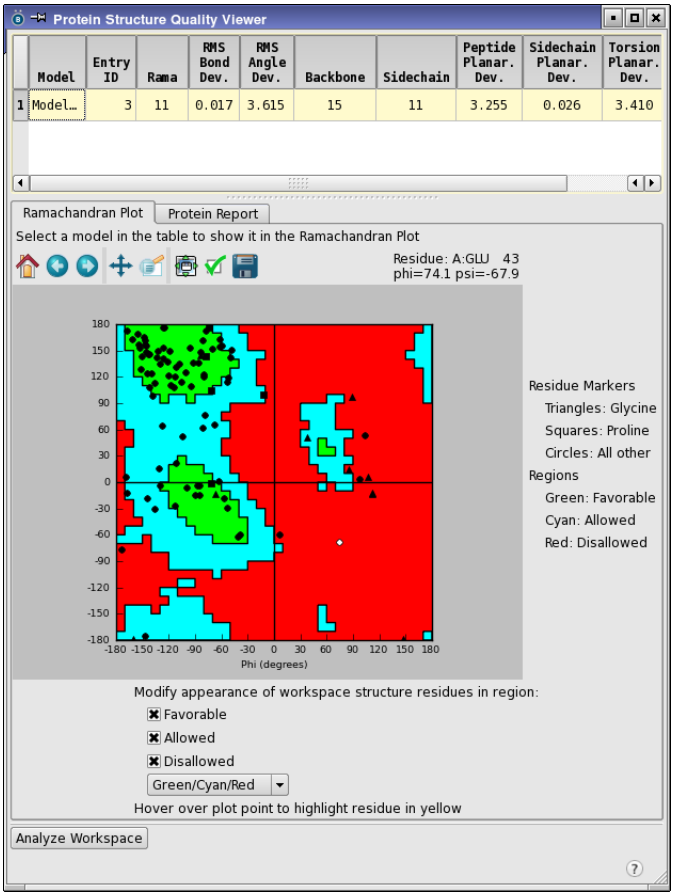

Figure 3.1. The Protein Structure Quality Viewer panel, Ramachandran Plot tab.

revert the color scheme to the original scheme, so you must change the scheme manually to revert it.

3.2 Protein Report

The property table displays a list of values of the property chosen from the Display menu. Selecting table rows zooms the Workspace view in to the structural features listed in those rows, highlights them with a change in representation, and selects them. The average of the selected property over the entire structure is displayed below the table.

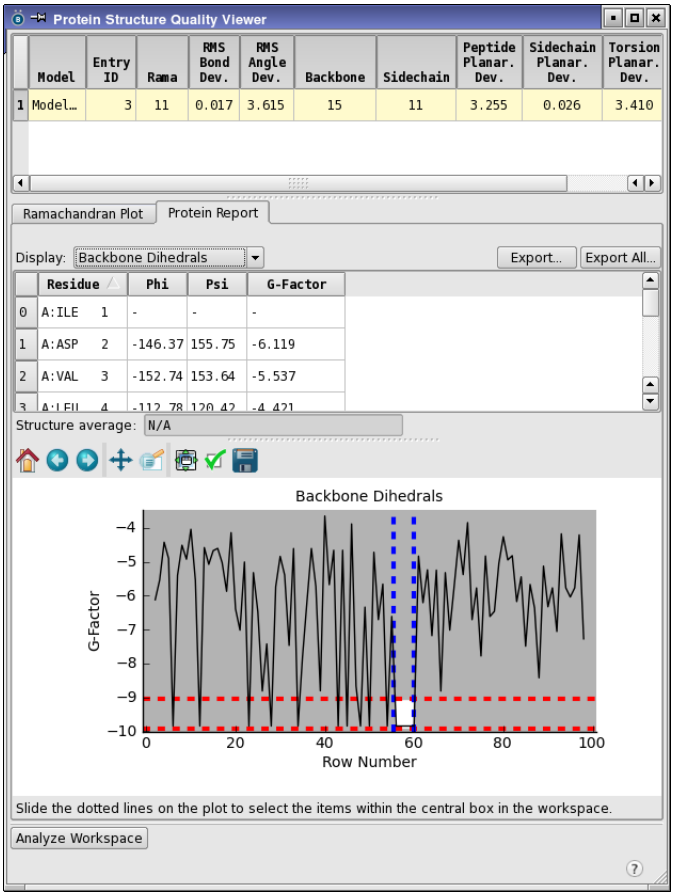

Figure 3.2. The Protein Structure Quality Viewer panel, Protein Report tab.

The property is plotted as a function of the row number in the table in the area below the table. The plot area has a toolbar, which is described in Section 3.3. The red dashed horizontal lines and blue dashed vertical lines can be dragged to highlight a portion of the plot of interest, which is given a white background, and the rest is gray. The atoms associated with the highlighted portion of the plot are interactively selected in the Workspace, and the rows for the points in the highlighted region are selected in the table.

If you want to export the values in the table, click Export or Export All. Export exports the current property table as a text file. Export All exports all property tables as a text file. Both buttons open a file selector in which you can navigate to the location and name the file.

3.3 Plot Toolbar

Both tabs have graphical displays, for which a toolbar provides some tools for manipulation of the plot and for saving an image of the plot. This is a generic toolbar, and some of the actions may not be useful in the current context. The panel has a toolbar that you can use to configure the plot or to save an image of the plot. The toolbar buttons are described below.

Reset
Reset the plot to the original pan and zoom settings.

Back
Display the previous view of the plot in the view history

Next
Display the next view of the plot in the view history

Pan/zoom
Pan the plot by dragging with the left mouse button, zoom by dragging with the right mouse button.

Zoom to rectangle
Drag out a rectangle on the plot to zoom in to that rectangle.

Configure subplots
Configure the margins and spacing of each plot in the panel.

Edit axis and curve parameters
Make settings for the title, range, labeling, and scale of the axes; the color, style, and width of lines; and the color, style, and size of markers.

Save image
Save an image of the plot to file. Opens a file selector in which you can browse to a location, select the image format, and name the image.

Analyzing Residue Properties

Identifying stable or unstable residues, or residues with desirable or undesirable properties may be a useful precursor to mutation studies. The Residue Analysis panel analyzes a protein to produce properties of the residues, including hydropathy, various energies, solvent-accessible surface areas, and rotatable bonds. These properties can be used to identify residues with either desirable properties or undesirable properties, which may suggest residues that could be mutated to improve the protein properties.

To open the Residue Analysis panel choose Tasks → Residue Analysis in the main window.

Before analyzing a protein, you should prepare it with the Protein Preparation Wizard. Calculating the energetic properties requires an all-atom structure with bond assignments. To analyze the properties of a protein, first display it in the Workspace, and then click Analyze Workspace. A job is run to calculate energetic properties. This job can take several minutes, and progress is reported in a bar at the bottom of the panel.

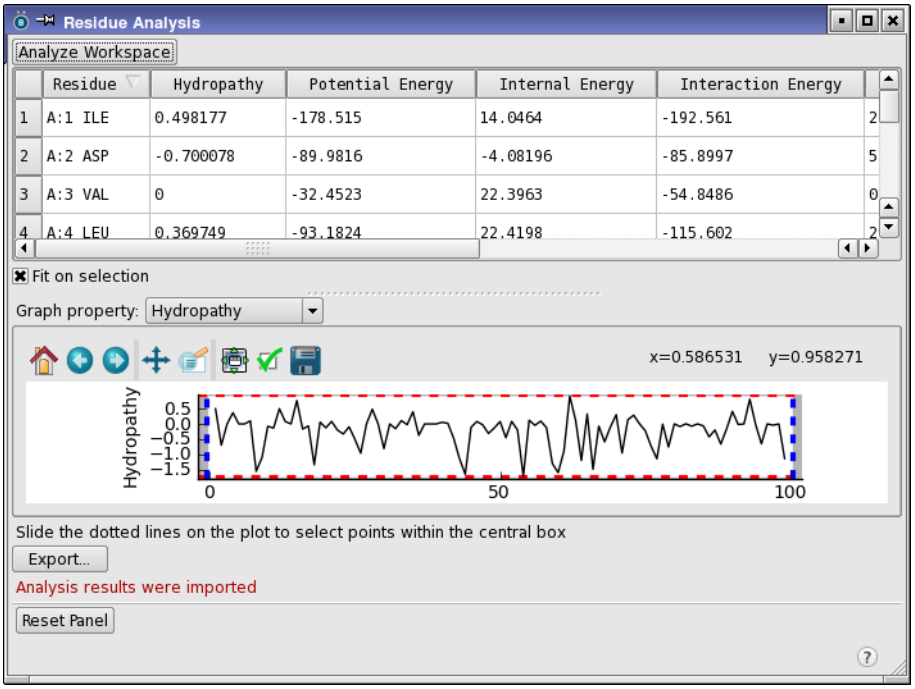

Figure 4.1. The Residue Analysis panel, with results.

When the job finishes, the table is filled in and the first property is plotted below. You can sort the table by clicking on the heading of the column you want to sort by. A second click changes the sort direction. The table columns are described in Table 4.1.

Table 4.1. Property columns in the Residue Analysis panel.

Column	Description
Residue	Residue identity: chain, residue number and insertion code, 3-letter name.
Hydropathy	Hydropathy calculated using the Kyte-Doolittle scale [6], normalized by the solvent accessible surface area.
Potential Energy	Sum of the residue-based internal energy and the non-bonded interaction energy (vdW, electrostatic) between the residue and the remainder of the system.
Internal Energy	Sum of energies arising from intra-residue bonded interactions (bonds, angles, torsions) and intra-residue non-bonded interactions (vdW, electrostatic).
Interaction Energy	Energy of interaction between this residue and all other atoms.
SASA (Non-polar)	Solvent-accessible surface area of nonpolar atoms of this residue
SASA (Polar)	Solvent-accessible surface area of polar atoms of this residue
SASA	Total solvent-accessible surface area of this residue
Rotatable bonds	Number of rotatable bonds in this residue

You can export the table data to a CSV file by clicking Export and then providing the file name in the file selector that opens.

To highlight one or more residues in the Workspace, select the table rows. If you have Fit on selection selected, the view zooms in (or out) so that these residues occupy most of the Workspace.

To examine a particular property for all residues, you can make a plot of the property as a function of residue position. Choose the property from the Graph property option menu to display the plot. You can use the plot to select residues in the table and highlight them in the Workspace, by moving the dotted lines to enclose the residues you are interested in. For example, you might want to select residues that have significantly larger or significantly smaller values of the properties than the average, by moving the upper or the lower red dotted line to enclose just those data points.

The toolbar provides tools for manipulation of the plot and for saving an image of the plot. This is a generic toolbar, and some of the actions may not be useful in the current context. The

panel has a toolbar that you can use to configure the plot or to save an image of the plot. The toolbar buttons are described below.

Reset
Reset the plot to the original pan and zoom settings.

Back
Display the previous view of the plot in the view history

Next
Display the next view of the plot in the view history

Pan/zoom
Pan the plot by dragging with the left mouse button, zoom by dragging with the right mouse button.

Zoom to rectangle
Drag out a rectangle on the plot to zoom in to that rectangle.

Configure subplots
Configure the margins and spacing of each plot in the panel.

Edit axis and curve parameters
Make settings for the title, range, labeling, and scale of the axes; the color, style, and width of lines; and the color, style, and size of markers.

Save image
Save an image of the plot to file. Opens a file selector in which you can browse to a location, select the image format, and name the image.

If you want to analyze another protein, click Reset to clear all the panel data.

Identifying Consensus Molecules

The Consensus Visualization panel helps you to identify conserved waters, counter ions, and ligands for a protein. To open this panel, choose Tools → Protein Consensus Viewer.

Homologs of the target protein are identified by a BLAST search. You can select a subset of these homologs to determine the consensus between them for the locations of waters, counter ions, and ligands. These homologs are aligned, both by sequence and by structure, to the target protein. The consensus between the positions of the waters, counter ions and ligands is then determined. Consensus analysis can help you to quickly identify moieties that are repeated among multiple structures, such as structurally important waters that should be included in modeling studies.

For a tutorial introduction to consensus visualization, see Chapter 5 of the *BioLuminate Quick Start Guide*.

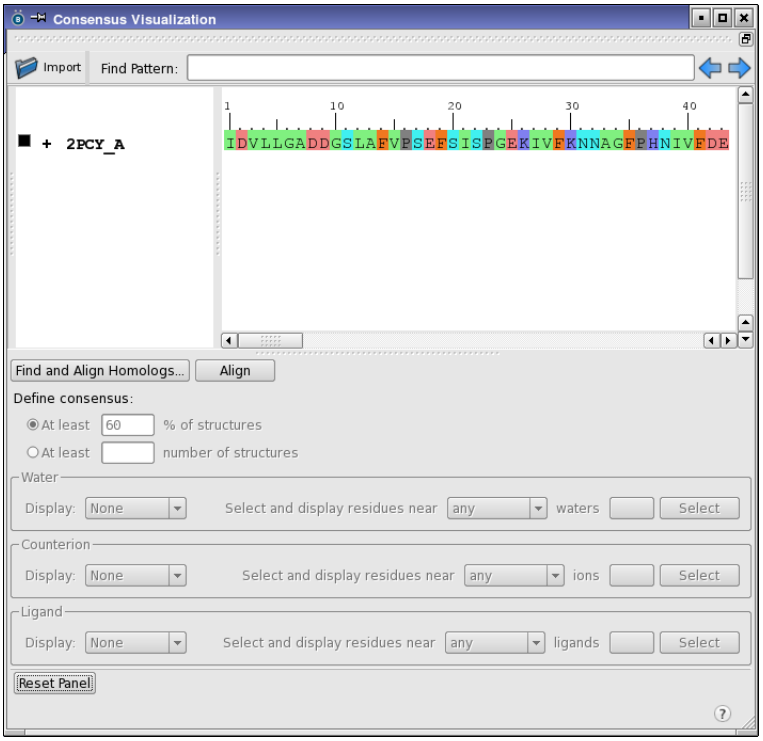

Figure 5.1. The Consensus Visualization panel after import.

5.1 Choosing the Target Protein

First, you must import your target protein. You can do this by clicking Import and choosing a source. The choices are:

- Browse for File—Open a file browser in which you can navigate to the desired location and select the file that contains the structure. The allowed file types are Maestro and PDB.
- From PDB ID—Import the structure from the specified PDB ID. Opens the Enter PDB ID dialog box, in which you can enter the PDB ID of the structure. The structure is retrieved from a local copy of the PDB if it is available, or from the RCSB web site, depending on the preference set for PDB retrieval.
- From Workspace—Import the structure that is displayed in the Workspace.

If you choose to prepare the protein beforehand in the Protein Preparation Wizard panel, you should ensure that you do not delete any of the molecules for which you are seeking a consensus. In particular, you might want to deselect Delete waters beyond *N* Å from het groups, or make the distance large enough to ensure that you have the relevant waters.

5.2 Finding and Aligning Homologs

The homologs that are used to identify consensus molecules can be imported, or you can run a BLAST search to locate homologs.

If you already have a set of homologs, you can simply import them with the Import button. Once they are imported you can align them structurally by clicking Align.

To run a BLAST search for homologs, click Find and Align Homologs. First, the Blast Search Settings panel opens. You do not usually need to change the BLAST search settings. When you are satisfied with the settings, click Start Job. A Job Progress dialog box replaces the Blast Search Settings dialog box, and displays the log file from the BLAST search.

After a few minutes, the job finishes and the BLAST Search Results dialog box opens, with the results of the search. The top ten results are selected in the table by default. You can change the number of top rows to select by entering the number of rows in the text box below the table, and clicking Select Top. You can also manually select rows in the table.

When you have finished selecting rows, click Incorporate Selected Rows. If you do not have a local installation of the BLAST or PDB databases, the search is done on the web, and a warning is displayed: “Multiple Sequence Viewer is attempting to access a remote server. Would you like to continue?” You can select Do not ask this question again, to prevent it from opening each time a structure is downloaded, then click OK.

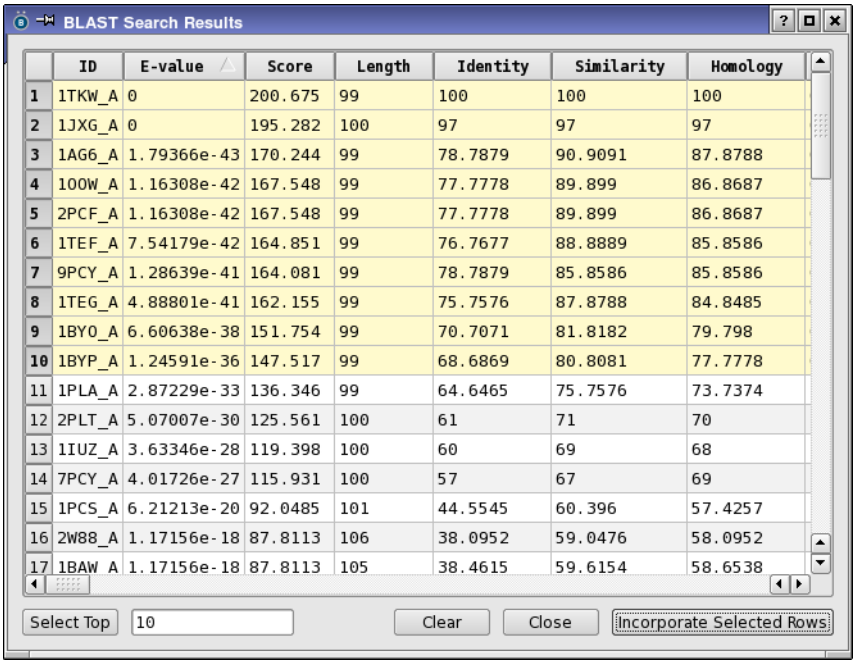

	ID	E-value	Score	Length	Identity	Similarity	Homology
1	1TKW_A	0	200.675	99	100	100	100
2	1JXG_A	0	195.282	100	97	97	97
3	1AG6_A	1.79366e-43	170.244	99	78.7879	90.9091	87.8788
4	1OOW_A	1.16308e-42	167.548	99	77.7778	89.899	86.8687
5	2PCF_A	1.16308e-42	167.548	99	77.7778	89.899	86.8687
6	1TEF_A	7.54179e-42	164.851	99	76.7677	88.8889	85.8586
7	9PCY_A	1.28639e-41	164.081	99	78.7879	85.8586	85.8586
8	1TEG_A	4.88801e-41	162.155	99	75.7576	87.8788	84.8485
9	1BYO_A	6.60638e-38	151.754	99	70.7071	81.8182	79.798
10	1BYP_A	1.24591e-36	147.517	99	68.6869	80.8081	77.7778
11	1PLA_A	2.87229e-33	136.346	99	64.6465	75.7576	73.7374
12	2PLT_A	5.07007e-30	125.561	100	61	71	70
13	1IUZ_A	3.63346e-28	119.398	100	60	69	68
14	7PCY_A	4.01726e-27	115.931	100	57	67	69
15	1PCS_A	6.21213e-20	92.0485	101	44.5545	60.396	57.4257
16	2W88_A	1.17156e-18	87.8113	106	38.0952	59.0476	58.0952
17	1BAW_A	1.17156e-18	87.8113	105	38.4615	59.6154	58.6538

Figure 5.2. The BLAST Search Results dialog box.

If an information box opens stating that problems were found when importing a structure, you can select Do not show this dialog again to prevent it from opening for each structure that has problems, and click OK. The structures are imported without any preprocessing, so they might have structural defects. For the purposes of this panel, it is generally acceptable to use structures from the PDB that have structural issues.

The homologs you selected are aligned, added to the sequence viewer, and displayed in the Workspace. All atoms are marked in all of the homologs.

5.3 Viewing Consensus Molecules

After the results are available, the next task is to decide how to define a consensus between the set of structures. There are two choices, available under Define consensus: a minimum percentage or a minimum number of structures. A match between the parent and a homolog for a particular molecule (such as a water) is obtained when any atom in the molecule in a homolog is within 2 Å of any atom in the same type of molecule in the reference (parent) structure. Consensus occurs when a match is found for the specified number or percentage of homologs.

For each of the three types of molecules, you can perform the following actions:

- Choose whether to display all, only the consensus, or none of the molecule of the given type, from the Display option menu. For example, viewing all the molecules gives an indication of whether a consensus exists, whereas viewing the consensus shows whether there is a strong enough consensus to consider the molecule as conserved.
- Select and display residues near the molecules of the given type. You can choose to display residues near any of the molecules or only the consensus molecules. The action is not performed until you click Select.
- Change the color of the residues that are near the molecules of the given type, by clicking on the color button, and choosing a color in the color selector that opens.

When displaying the molecules, the identity of any consensus molecule can be ascertained by moving the cursor over the structure in the Workspace and viewing the text in the status bar, below the Workspace. Consensus waters and ions are displayed as spheres, consensus ligands as ball-and-stick, and they are highlighted with a silhouette (which can be changed with Style → Highlights, Text and Arrows).

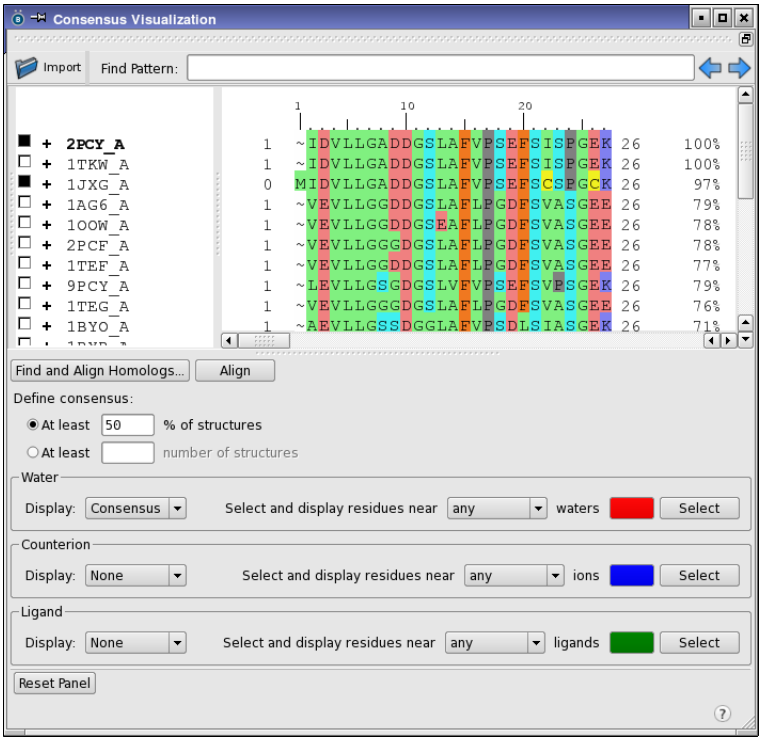

Figure 5.3. The Consensus Visualization panel with results.

Identifying Reactive Residues

You may need to identify reactive residues in a protein, so that you can mutate them to improve the protein properties. This can be done in the Reactive Protein Residues panel, which you open by choosing Tools → Reactive Residue Identification.

Reactive residues are identified by matching residue patterns in the sequence. Four patterns are provided by default, for the common reactions: deamidation, oxidation, glycosylation, and proteolysis. You can use these patterns or you can set up and use your own patterns.

To identify reactive residues, include the protein you want to analyze in the Workspace, and click Analyze Workspace. The structure in the Workspace is analyzed to identify residues that match the patterns.

The results are listed in the table, which shows the reaction type, the reactive residues identified, the solvent-accessible surface area of the reactive residues, their percentage exposure to

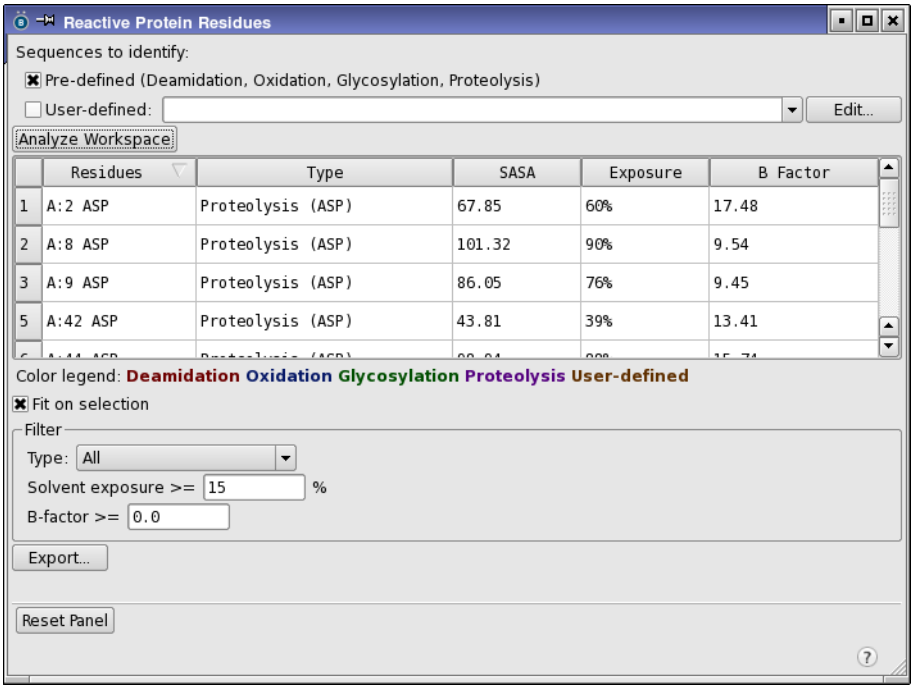

Figure 6.1. The Reactive Protein Residues panel.

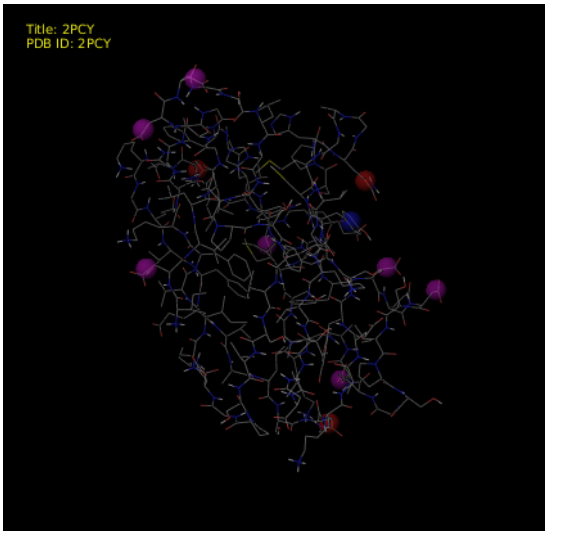

Figure 6.2. Reactive residue sites in 2PCY.

solvent, and their B-factors (if available). The B-factor shown is the average over all atoms in the residue.

You can use the Show option menu to show only the results for a particular reaction type. You can also apply filters on the percentage solvent exposure and the B-factor.

To sort the table by the values in a column, click on the column heading. Click again to change the direction of the sort. The column by which the table is sorted is indicated by an arrow on the right side of the heading, which also indicates the sort direction.

The reactive sites are marked with spheres in the Workspace. The spheres are colored according to the reaction type, and the color legend for the spheres is given below the table. When you select a table row, the residue is highlighted in the Workspace. If you have Fit on selection selected, the view zooms in to that residue

If you want to define your own reactive groups, click Edit, to open the Edit Patterns dialog box. This dialog box lists the patterns in a table, giving the pattern name, the definition, and a hotspot index. To edit an existing pattern, double-click the table cell that you want to edit, and enter the changes. The lower part of the panel explains the syntax for the patterns in the Definition column. The syntax is an extended PROSITE syntax, which allows you to specify secondary structure and some properties:

- Standard IUPAC one-letter (upper case) codes are used for all amino acids.
- Lower case x is used for any amino acid.

- Each element of a pattern is separated with a - symbol.
- Residues that are permitted at a given position are listed between square brackets, e.g. [ACT] means one of Ala, Cys, or Thr, or in other words, only Ala, Cys, or Thr can appear at this position.
- Residues that are not permitted at a given position are listed between curly brackets, e.g. {GP} means not Gly and not Pro, or in other words, any residue but Gly or Pro can appear at this position.
- Repetition is indicated using parentheses, e.g. A(3) means Ala-Ala-Ala, G(2,4) means between 2 to 4 Gly residues.
- The following lower case characters can be used for residue types:
  - a—acidic residue: [DE]
  - b—basic residue: [KR]
  - o—hydrophobic residue: [ACFILPWVY]
  - p—aromatic residue: [WYF]
- The following lower case characters can be used to restrict residue types by property:
  - s—solvent-exposed residue
  - h—residue in helical region
  - e—residue in extended (beta strand) region
  - f—flexible residue, defined as having a B-factor above the average over all residues

  These four characters can be appended to a residue type to restrict the type of residue, e.g. Ah means Ala in a helical region.

Some examples of valid and invalid patterns are given below, with comments.

N-{P}-[ST]	Asn-X-Ser or Asn-X-Thr, X is not Pro
N[fs]-{P}[fs]-[ST][fs]	as above, but all residues flexible or solvent exposed
Nfs-{P}fs-[ST]fs	as above, but all residues flexible and solvent exposed
Ns{f}	Asn, solvent exposed and not in flexible region
N[s{f}]	Asn, solvent exposed or not in flexible region
[ab]{K}f{s}	acidic OR basic, except for flexible and non-solvent-exposed Lys
Ahe	Ala, helical and extended - no match is possible
A[he]	Ala, helical or extended
A{he}	Ala, not helical or extended
[ST]	Ser or Thr
ST	Ser and Thr - no match possible

The hotspot index is the index of the reactive residue in the pattern.

Predicting Aggregation Regions

Protein aggregation often occurs via hydrophobic regions. You can locate and analyze these regions using the Aggregation Surface panel, which you open by choosing Tasks → Aggregation Surface. Aggregation regions are defined by identifying clusters of exposed hydrophobic residues, and creating a molecular surface that is colored red in the regions near these residues.

Once you have located potential aggregation regions on a protein, you might want to mutate residues in these regions to reduce the tendency to aggregate.

7.1 Creating an Aggregation Surface

Follow the steps below to create a surface that displays the likely aggregation regions.

1. (Optional) Select atoms to define the structure for the surface.

   If you want to create a surface for the entire protein and you have only one entry in the Workspace, you can skip this step. Otherwise, you can select atoms to create a surface for an entire entry, a chain, or the selected atoms with their hydrogens.

2. Choose an option for the part of the Workspace structure that you want to create the surface for.

   There are several options. If no atoms are selected in the Workspace, only the first is available.

   - Workspace—Compute the surface for the structure in the Workspace. You must have only a single entry in the Workspace.
   - Entire entry containing selected atoms—Compute the surface for the entry that contains the selected atoms. This is useful if you want to view how the aggregation surface on one protein matches residues in another protein.
   - Entire chain containing selected atoms—Compute the surface for the entire chains that contain the selected atoms. This is useful for computing a surface for one chain of a multi-chain protein, for example.
   - Selected atoms and attached hydrogen atoms only—Compute the surface for the selected atoms with their attached hydrogens only. This option allows you to calculate the surface for only part of a protein chain, for example.

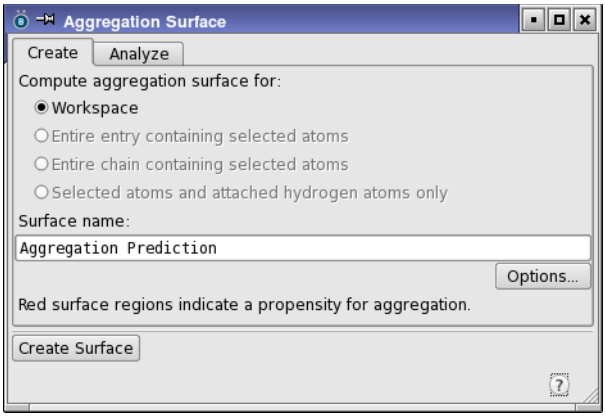

Figure 7.1. The Create tab of the Aggregation Surface panel.

3. Enter a name in the Surface name text box, if you want to change the default name.

   This name is used in the Manage Surfaces panel (Style → Create and Manage Surfaces → Surface Manager) to identify the surface.

4. Click Create Surface.

   A progress bar is displayed above the buttons while the surface is being created. Surface creation should take less than a minute.

When the surface is created, it is displayed in the Workspace and other surfaces are hidden. To display more than one surface, use the Manage Surfaces panel (Window → Show → Surface Manager).

The aggregation surface is a molecular surface that is created for a large probe molecule, representing part of a protein. It is colored red in the regions near the hydrophobic residue clusters. The side chains of these residues and the alpha carbons are also colored red, and the remaining backbone atoms are colored gray. The side chains are displayed in ball-and-stick representation, and the backbone is displayed as lines. Residues that are in contact with these residues are displayed as lines in gray. All other residues are hidden, so you see only the residues that contribute to the aggregation regions and very near neighbors. The surface is semi-transparent, so you can see the atoms inside the surface. The residues involved in the clusters are colored red in the Workspace sequence viewer.

The surface is stored with the project entry, just like any other surface. If you want, you can create more than one surface for the Workspace contents. For example, you might have two proteins displayed and want to analyze the aggregation regions at the interface.

7.2 Setting Options for Aggregation Surfaces

You can control some of the parameters of the quality and appearance of the aggregation surface and the parameters that are used for locating aggregation regions in the Aggregation Surface - Options dialog box. To open this dialog box, click Options in the Aggregation Surface panel. The parameters you can set are:

- Surface grid spacing—Set the grid spacing in angstroms for generation of the surface. A smaller number results in a smoother surface, but takes longer to generate the surface.
- Probe radius—Set the radius of the probe for defining the surface. This is the radius of the sphere that is rolled over the van der Waals surface to create a Connolly surface. The large default radius is intended to model a protein probe. Hydrogens are included when creating the surface. See Section 12.1.2 of the *Maestro User Manual* for details on the construction of the surface.
- Transparency—Set the default transparency of the surface. The transparency can be changed in the Surface Display Options dialog box—see Section 12.4.2 of the *Maestro User Manual*.
- Radius to find neighbors—Specify the distance cutoff for finding hydrophobic neighbors to identify aggregation regions. The distance between any side-chain heavy atom in one hydrophobic residue and any side-chain heavy atom in another hydrophobic residue must be less than this cutoff for the residues to be counted as neighbors.
- Hydrophobic neighbors required for site—Specify the minimum number of hydrophobic neighbors required to include a residue (site) in an aggregation region.
- Buried residue SASA—Specify the maximum solvent-accessible surface area (SASA) for a residue to be regarded as buried. Buried residues are not included in aggregation regions. The SASA is measured for the heavy atoms only: it does not include hydrogens.

7.3 Analyzing the Surface

The Analysis tab provides some tools for analyzing which residues contribute to the aggregation regions. To analyze a surface, choose it from the Surface option menu and click Analyze.

When the analysis finishes, the table is filled in with a list of residues that contribute to the aggregation regions of the surface. For each residue, its contribution to the surface and the index of the group to which it belongs is listed in the table. The contribution is a count of surface elements, which is roughly proportional to the surface area due to that residue. A group is a set of residues that contribute to the same aggregation region, defined by the Radius to find neighbors setting in the Aggregation Surface - Options dialog box.

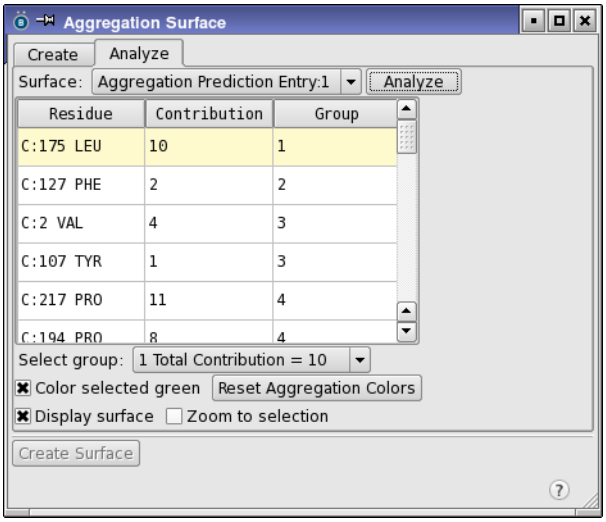

Figure 7.2. The Analyze tab of the Aggregation Surface panel.

When you select a row in the table, the residue is selected in the Workspace. You can select entire aggregation regions (groups) by choosing the group from the Select group option menu. The residues are selected in the table and the Workspace. The option menu shows the sum of contributions to the group as well as the group index.

If Color selected green is selected, the red surface patch associated with the selected residues is changed to green, and the residues themselves are colored green also. You can reset the colors to the original aggregation color scheme (red and gray) by clicking Reset Aggregation Colors.

If you want to zoom in to the selection in the Workspace when you select residues or groups, select Zoom to selection.

You can display or hide the surface by selecting or deselecting Display surface. As the color scheme is applied to both the residues and the surface, the residue coloring still changes with selection when the surface is hidden.

7.4 Using the Results for Mutation Studies

To reduce the likelihood that the protein will aggregate, you might want to mutate residues in the aggregation regions.

One way of doing this is to perform a set of mutations to reduce the size of the aggregation regions. You can create a set of structures that have single mutations at selected sites, which you choose from the residues in the aggregation regions, as follows:

1. Right-click in the Workspace (not on an atom) and choose Visible → Select.

   The visible atoms, which include all the aggregation residues, are selected, along with the neighbors. You can also use the table in the Analyze tab to select residues or groups rather than the entire set of residues.

2. Choose Tasks → Residue Scanning → Perform Calculation to open the Residue Scanning panel.

3. Select Analyze only selected Workspace residues.

4. Click Analyze Workspace. You may be prompted to choose a chain, as only one chain at a time is mutated.

The residues that were found to contribute to the aggregation region are listed in the table in the Residues tab. You can then select any of them for mutation, define the mutations, and run the job. See Chapter 15 for details of setting up a residue scanning job.

Another option is to mutate a single residue or a loop. Mutating a loop ("loop swap") is useful if you want to mutate more than one residue and the residues are adjacent. For this purpose, you can use the Residue and Loop Mutation panel. The loop to change is defined by selecting the residues in the Workspace. You can take advantage of the fact that you only have the aggregation residues and their neighbors visible in the Workspace to choose the residues for the loop to swap, or you can use the Analysis tab to select the residues, and use the selection. See Chapter 13 for details of using this panel.

Analyzing Protein Interface Interactions

It can be useful to analyze the specific nature of the interactions at the interface of two proteins. BioLuminate provides this capability in the Protein Interaction Analysis panel. The analysis locates residues in one protein that are within a given distance of residues in another protein, and presents counts of hydrogen bonds, salt bridges, disulfide bonds, pi-pi stacking interactions, and van der Waals clashes, and reports the van der Waals surface complementarity and buried solvent-accessible surface area.

To open the Protein Interaction Analysis panel, choose Tasks → Protein Interaction Analysis in the main window.

The proteins whose interface you want to analyze must be part of a single project entry. The could come from a protein-protein docking run, for example, or an antibody-antigen complex, or a multi-chain protein from the PDB. Once you have a structure, from whatever source, it must be properly prepared, for example in the Protein Preparation Wizard panel.

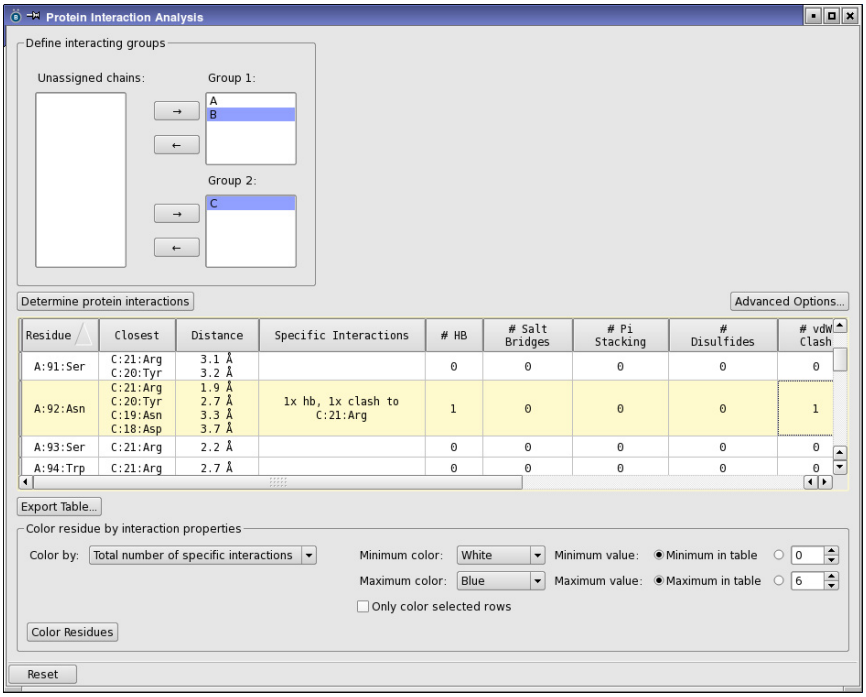

Figure 8.1. The Protein Interaction Analysis panel.

8.1 Selecting the Interacting Proteins

The two proteins whose interface you want to analyze must be identified in the Workspace structure. This is done by defining two groups of chains, one for each protein. When you open the Protein Interaction Analysis panel, the chains are listed in the Unassigned chains list in the Define interacting groups section.

To define the two proteins:

1. Select chains for the first group in the Unassigned chains list.

   You can select multiple chains with the usual shift-click and control-click actions.

2. Click the right arrow button (→) for Group 1.

   The chains are listed in the Group 1 list and removed from the Unassigned chains list.

3. Repeat the above two steps for Group 2.

To move chains between Group 1 and Group 2:

1. Select the chains in the Group 1 list or the Group 2 list.
2. Click the left arrow button (←) for the group to unassign them.

   The chains are moved to the Unassigned chains list, and remain selected.

3. Click the right arrow button (→) for the other group to reassign them.

You do not have to include all the chains in the analysis, provided that there is one chain in each group. Chains that are unassigned are not analyzed.

8.2 Running the Analysis

Once you have defined the two groups, you can analyze the interactions at the interface, by clicking Determine Protein Interactions. This may take a minute or two. When the analysis is done, the results are displayed in the table.

The analysis is performed using settings that characterize the interactions, such as interatomic distances and angles. You can change the settings in the Advanced Options dialog box, which you open by clicking Advanced Options. These settings affect the values that are listed in the table after the analysis is done. The settings are listed below. To return to the default settings, click Reset.

List closest interaction neighbors within

Set the cutoff for determining the closest interaction neighbors for the interaction analysis. Any residue that has any atom within the specified distance of the target residue is considered a neighbor of that residue. The default is 4.0 Å.

Ignore interactions between backbone atoms

Do not include backbone–backbone interactions in the interaction analysis.

Hydrogen bonds

Specify the criteria for determining whether a hydrogen bond exists. The four atoms involved in the hydrogen bond are designated D–H...A–X, where D is the donor atom and A is the acceptor atom. The default values are the Maestro defaults.

- Minimum acceptor angle—Set the minimum acceptor angle H...A–X, in degrees. The default is 90°.
- Minimum donor angle—Set the minimum donor angle D–H...A, in degrees. The default is 120°.
- Maximum distance—Set the maximum H...A distance, in angstroms. The default is 2.5 Å.

Salt bridges

Specify the maximum distance between an ion and a protein atom for detecting a salt bridge, in the Maximum distance box. The default is 4.0 Å.

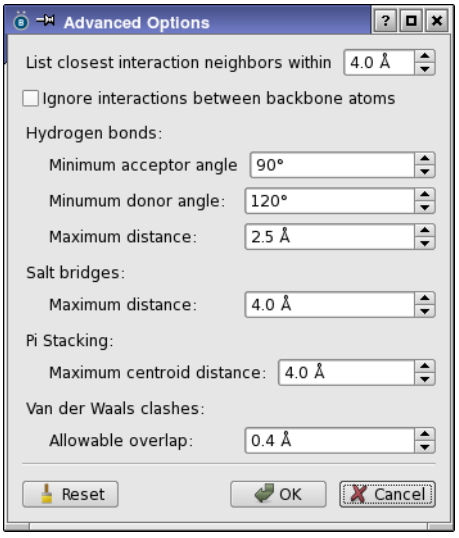

Figure 8.2. The Advanced Options dialog box.

Pi stacking

Set the maximum distance between the centroids of the two aromatic rings in the Maximum centroid distance box. The default is 4.0 Å.

Van der Waals clashes

Set the maximum overlap distance of the van der Waals spheres of any two atoms in the Allowable overlap box. If $R_A+R_B-R_{AB}$ is greater than the allowable overlap, the atoms are considered to clash, where R_A and R_B are the van der Waals radii, and R_{AB} is the distance between atoms A and B. The default is 0.4 Å.

8.3 Examining the Results

The results of the analysis are listed in the table in the center part of the panel. The table columns are described in Table 8.1. You can click on a column heading to sort the table by the values in that column (except for the Specific Interactions column). When you select table rows, the residues for those rows are selected in the Workspace, and the view zooms in to those residues.

Table 8.1. Interaction table columns.

Column	Description
Residue	This column lists residues in either group that have contact-type interactions with residues in the other group.
Closest	This column lists residues from the other group that are within a specified distance of the residue listed in the Residue column. The default distance is 4.0 Å.
Specific Interactions	Text list of specific interactions between the residue listed in the Residue column and residues in the other group. The list covers hydrogen bonds, salt bridges, pi-pi interactions, disulfides, and van der Waals clashes.
# HB	Number of hydrogen bonds between the residue listed in the Residue column and residues in the other group. The criteria for detecting hydrogen bonds can be changed in the Advanced Settings dialog box, prior to the analysis.
# Salt Bridges	Number of salt bridges between the residue listed in the Residue column and residues in the other group.
# Pi Stacking	Number of pi-pi stacking interactions between the residue listed in the Residue column and residues in the other group.
# Disulfides	Number of disulfide bonds between the residue listed in the Residue column and residues in the other group.

Table 8.1. Interaction table columns. (Continued)

Column	Description
# vdW Clash	Number of van der Waals clashes between the residue listed in the Residue column and residues in the other group. A clash is defined as an overlap of the van der Waals radii of two atoms by more than a specified cutoff. The default is 0.4 Å.
vdW Complementarity	Van der Waals shape complementarity between the residue listed in the Residue column and residues in the other group, as defined in Ref. 10.
Buried SASA	Fraction of the solvent-accessible surface area of the residue listed in the Residue column that is buried by the interaction with residues in the other group.

If you want to import the results into another application, you can export them to a CSV file by clicking the Export Table button, and naming the file in the file selector that opens. The file is comma-separated with a heading row. Table cells that have multiple rows are exported as double-quoted text with line breaks embedded.

8.4 Viewing Interaction Properties in the Workspace

You can color residues in the Workspace by the value of one of three interaction properties. The color scheme is a linear ramp between the color for a chosen minimum value and the color for a chosen maximum value of the property. The minimum and maximum values can be set to the range of values in the table or to fixed values, which allows comparisons of proteins on the same scale.

You can choose the property from the Color by option menu. The choices are:

- Total number of specific interactions—Total number of interactions described in the Specific Interactions column (sum of the values in the # HB, # Salt Bridges, # Pi Stacking, # Disulfides, and # vdW Clash columns).
- vdW Complementarity—van der Waals shape complementarity, as listed in the table.
- SASA—Buried solvent-accessible surface area, as listed in the table.

When you choose an option from this menu, the text boxes for the minimum and maximum values are updated to reflect the possible range of values.

The colors that represent the chosen minimum value and maximum value of the interaction are chosen from the Minimum color and Maximum color option menus. The minimum and maximum values for the color display are set with the value options, described below. The color for the minimum is applied to any value below the chosen minimum for display; likewise the color for the maximum is applied to any value above the chosen maximum for display. The colors between the two limits are obtained by a linear interpolation of the RGB color values.

The minimum value and maximum value for the color scheme can be chosen by selecting one of the Minimum value and Maximum value options. There are two options for each value:

- Minimum/Maximum in table—Set the values to use for coloring to the minimum value and maximum value found in the table. The color scheme so defined is then relative to the range observed, rather than fixed.
- Supplied value—Set the values to use for coloring to the values given in the boxes. The default minimum value is zero; the default maximum is the allowed maximum for the total number of specific interactions, 1.0 for the vdW complementarity, and 100% for the buried SASA.

If you only want to color some of the residues, you can select the table rows for the residues you want to color, and select Only color selected rows.

When you have chosen the property, set up the color scheme, and optionally chosen residues to color, click Color Residues to apply the color scheme.

Locating Large-Scale Motions

Finding large-scale motions in proteins can provide information on the flexibility or rigidity of parts of the protein, on domain movements, or give some clues to biochemical processes. Trajectories from molecular dynamics simulations can show the large-scale motions of proteins, but these simulations are not resolved into individual modes, and they require a large amount of computational resources. The lowest vibrational modes of a protein can be determined and visualized using the Low Mode Vibrational Sampling panel. You can open this panel by choosing Tasks → Low Normal Mode Analysis → Calculate or Visualize.

Before running the calculation, you should ensure that the protein is properly prepared, using the Protein Preparation Wizard (Tools → Protein Preparation). You should remove waters and solvent molecules so that you are analyzing just the protein (and its ligands, if any).

First, the input structure is minimized (using the PRCG method for a maximum of 10000 iterations to a gradient convergence threshold of 0.05 kJ $mol^{-1}Å^{-1}$). This ensures that the structure is at its minimum, which is important for generating the vibrational modes. The vibrational modes are generated as a set of structures sampled at regular intervals along a full cycle of the vibrational mode. The rotational and translational modes (the trivial modes) are discarded. The vibrational modes are then visualized by displaying the structures in sequence as a "movie" in the Workspace.

To set up the calculation:

1. Select the number of vibrational modes you want to view in the Number of vibrational modes to view text box.
2. Set the number of frames to generate per vibrational cycle in the Number of frames per mode text box.

   Each frame is a snapshot of the structure at a particular point in the vibration between the classical limits. The full cycle of a vibration is divided evenly to define the coordinates used for the snapshots, so this number should be divisible by 4.
3. Set the maximum amplitude of vibration in the Vibrational amplitude text box.

   This is the maximum displacement of the fastest-moving atom. For proteins, the fastest moving atom could potentially move a large distance if the motion involves a long loop, for example. This choice is somewhat arbitrary: you might for example want to exaggerate the motions to make them easier to see.

4. Click Run.

A job is started that generates the series of structures for each vibration. The limit on the number of atoms that can be processed depends on the memory available on the machine, but with 4GB of memory, it is about 5000. For example, 1ETT, with about 4800 atoms, runs successfully. The job can take several hours to run, depending on the size of the protein.

To visualize the vibrational modes:

1. Import the results of the job by clicking Browse and selecting the `.com` file for the job in the file selector that opens.

   The results are added to the project as an entry group.

2. Specify the mode you want to view in the Vibrational mode to visualize box.
3. Specify the duration of each frame, in seconds.

   The speed at which the structures can be displayed depends on the size of the structure. For structures of 5000 atoms, an interval of 0.1 sec is possible and produces a reasonable animation.

4. Select Loop to loop continuously through the frames, so that the vibration goes through multiple cycles.
5. Use the play controls to start and stop the animation.

   You can rotate the Workspace during the animation to get a better view of the moving parts of the structure.

When you have finished viewing the modes, you should clean up the structures in the project, by clicking Remove Structures from Project, or deleting the entry group in the Project Table panel.

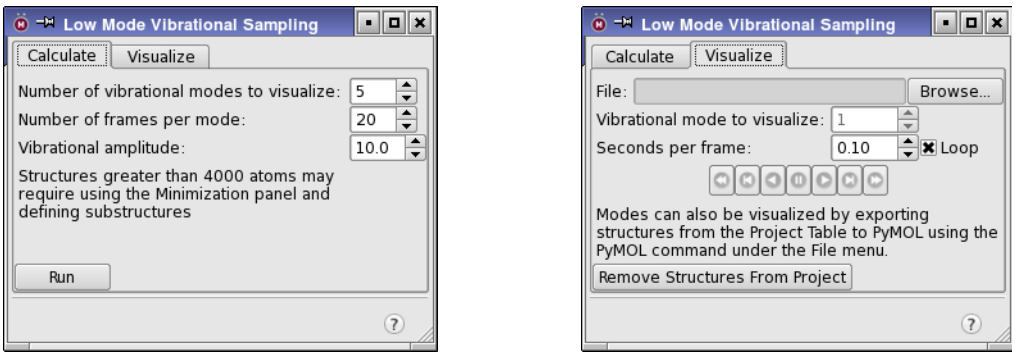

Figure 9.1. The Low Mode Vibrational Sampling panel.

Predicting Peptide Properties

10.1 Alpha Helix Stability

Determining the stability of alpha-helical peptides is important for their use as therapeutic agents. You can obtain an estimate of the stability by using the Peptide Helicity panel, which you can open by choosing Tasks → Peptide Alpha Helicity. The primary purpose of this panel is to provide information on the relative stabilities of a series of small peptides, although you can also obtain information on a single peptide.

The alpha helical tendency of one or more peptides is determined from a molecular dynamics simulation by tracking i to i+4 hydrogen bond formation, and other indications of helical structure. The prediction is made entirely from the sequences, by building them as idealized alpha-helices, and then performing a molecular dynamics simulation in water using simulated annealing, to simulate experimental melting experiments. At the end of the simulation, averages are taken to determine the values of properties that can be used to determine the helicity of the sequences.

The simulations are set up in the Start Simulations tab, and the results are presented in the Results tab.

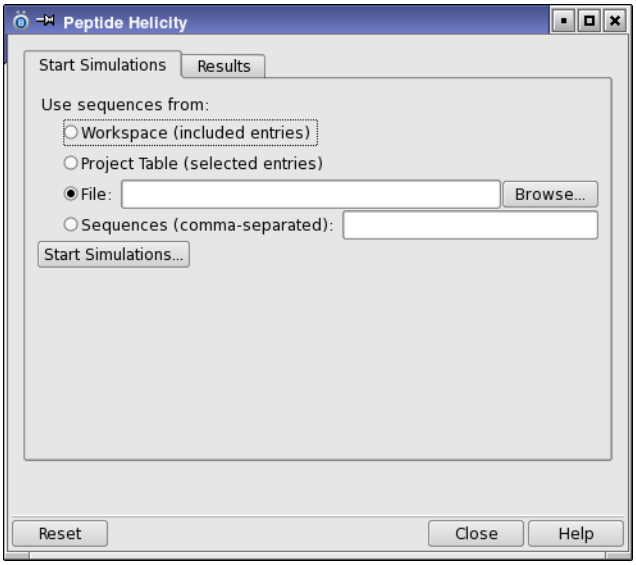

Figure 10.1. The Peptide Helicity panel.

To run a simulation:

1. Select an option for the sequences to simulate under Use sequences from.

   The options are Workspace (the included entries), Project Table (the selected entries), File, or Sequences. If you choose File, enter the file name in the text box or click Browse to navigate to and select the file. If you choose Sequences, type or paste the sequences into the text box, separated by commas. Each sequence that you provide is simulated separately in a single job.

2. Click Start Simulations.

   A Start dialog box opens, in which you can set the job name, select a host, and specify the number of processors. Click Start in this dialog box to submit the job to a host for execution.

   Since the simulations are likely to take many hours, it is a good idea to run the job on a multiprocessor host. The simulation uses a 3D domain decomposition of the simulation box, so you can specify the number of processors for each dimension in the decomposition (labeled x, y, and z). See Section 3.10 of the *Desmond User Manual* for more information. A good choice is 2 processors for each, on an 8-core CPU.

When the simulation finishes, you can review the results of the simulations in the Results tab. Click Load Output File read the results, which are in a CSV file. A file selector opens, in which you can navigate to and select the output CSV file. The results are presented in a table, whose columns are described in Table 10.1. The averages are taken over the length of the simulation.

If you want to run another calculation, click Reset to clear all panel data.

Table 10.1. Columns of the Results table in the Peptide Helicity panel.

Column	Description
Title	Structure title
Structure No	Structure number
α-Propensity	Average helical propensity, defined as the fraction of (i, i+4) backbone hydrogen bonds.
H-bonds	Average number of alpha helical H-bonds
Residues Count	Average number of residues involved in two H-bonds
Residues Fraction	Average fraction of residues satisfying the phi/psi criteria for helicity individually
Dihedrals	Average fraction of alpha-helical dihedrals, as a fraction of the maximum number present in the idealized helix.

10.2 QSAR from Sequences

Predicting the property of a peptide based on its sequence can be useful when designing and selecting peptides for application as therapeutics. You can build and apply a QSAR-type model that relates properties of peptide sequences to an observed property in the Peptide QSAR panel.

The independent variables in these models are sets of amino acid properties taken from the literature [7–9], and do not need to be explicitly provided. Only the observed property is needed, for the training set and the test set of the model.

To open the Peptide QSAR panel, choose Tasks → Peptide QSAR. The model is set up in the Setup tab, and the results are presented in the Results tab.

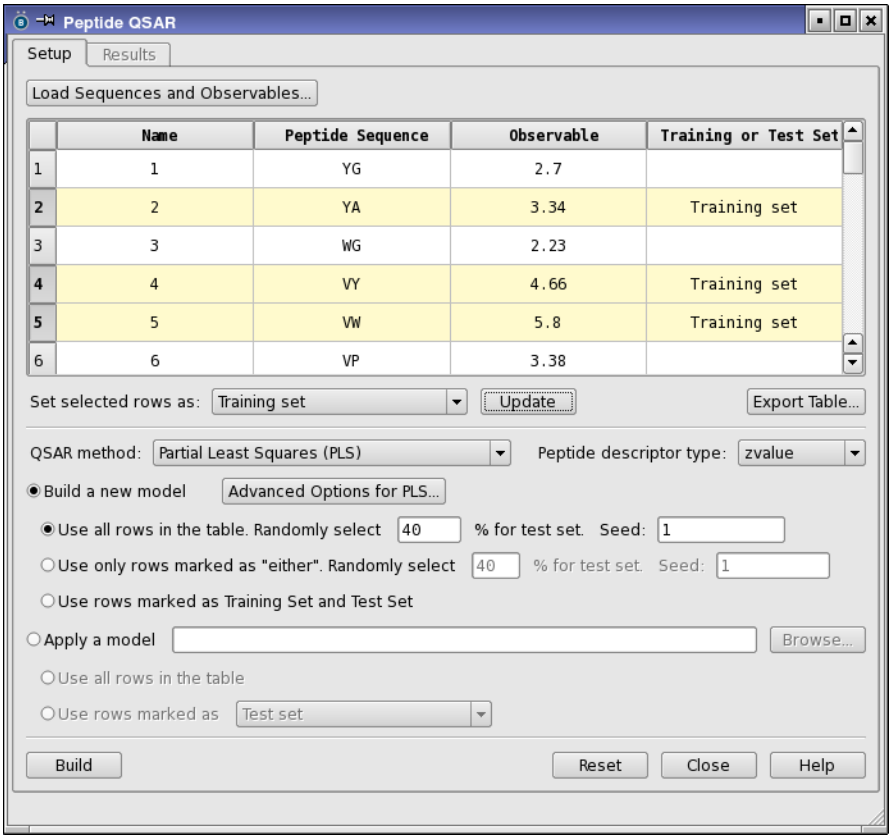

Figure 10.2. The Setup tab of the Peptide QSAR panel.

10.2.1 Loading the Sequences and Properties

The first task is to load the sequences and observable property for each sequence, whether you are building a model or applying a model. To do this, click Load Sequences and Observables, and make choices in the Load Sequences and Observables dialog box.

The sequences can come from peptide structures in the project, from a Fasta file, or from a CSV file, and must be all of the same length. You can choose the source of the sequences from the Load sequences from option menu.

- Project Table (selected entries)—Select the entries in the Project Table that contain the peptide sequences.
- Workspace (included entries)—Include the entries in the Workspace that contain the peptide sequences. The sequences are displayed in the Workspace sequence viewer.
- CSV file—Read the sequences and the observables from a CSV file.
- Fasta file—Read the sequences from a Fasta file.

The choice you make determines which controls are displayed in the rest of the dialog box for the selection of the observables.

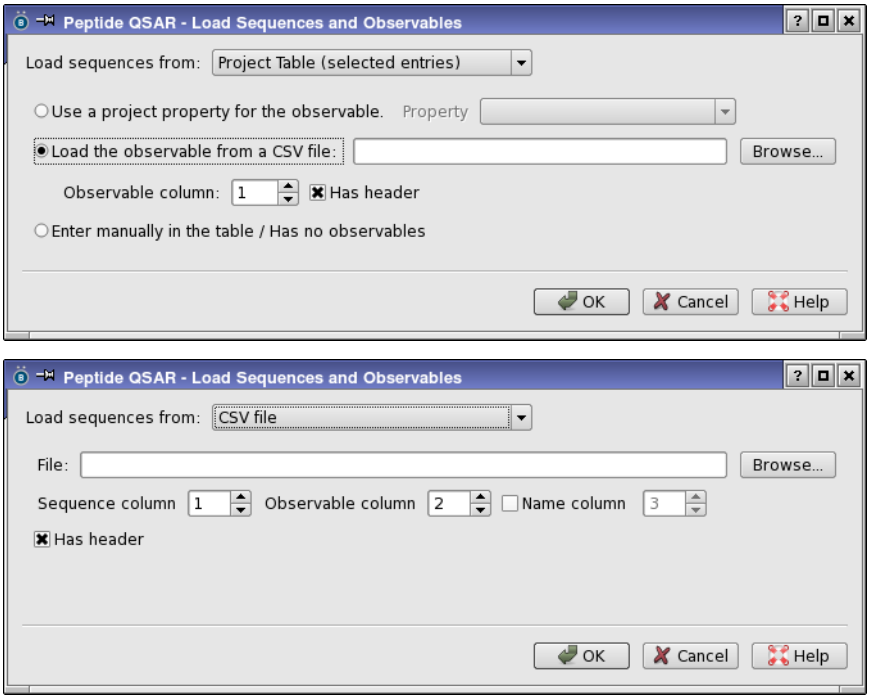

Figure 10.3. The Load Sequences and Observables dialog box

- If you load the sequences from a CSV file, the observable must also be in the CSV file. You can specify the file with the usual tools (File text box, Browse button), and you can specify the column that contains the sequence and the column that contains the observable. If the sequences have names, you can also check the box for the name and specify the column that contains the name. If the CSV file has a header row, select Has header.
- If you load the sequences from the Project Table or the Workspace, you can use a project property for the observable, and choose the property from the Property option menu.
- For any of the methods you can load the observable from a CSV file. You must specify the CSV file, the column in the file that contains the observable, and indicate whether the file has a header row.
- You can also leave the observable property undefined, by choosing Enter manually in the table / Has no observables. You can do this when building a model if you plan to enter the values manually in the table. If you are making predictions for new sequences, select this option to load sequences without observables. This option is not available when loading sequences from a CSV file.

Click OK when you have made your choices. The dialog box closes, and the sequence table in the Peptide QSAR panel is filled in (see Figure 10.2 on page 59). If you chose to leave the observables undefined, you can edit the table cells to supply the values.

10.2.2 Building a QSAR Model

When you have a set of sequences with an observable property, you can proceed to building a QSAR model. Select Build a new model to activate the model-building settings.

First, you need to choose a training set and a test set. There are three options:

- You can choose the sets explicitly. Select Use rows marked as Training Set and Test Set. Select the rows for one of the sets in the table, choose the set from the Set selected rows as option menu, and click Update. Do the same for the other set. You should of course only select rows that have observable values.
- You can use all the sequences, and assign the training and test sets randomly. Select Use all rows in the table, and set the percentage for the test set in the Randomly select text box. You can also specify a seed for the random number generator.
- You can use a subset of the sequences, and assign the training and test sets randomly. Select Use only rows marked as "either", and set the percentage for the test set in the Randomly select text box. Select the rows to use in the table, choose Either training or test from the Set selected rows as option menu, and click Update. You can also specify a seed for the random number generator.

Next, choose the type of amino acid descriptor set to be used as the X (independent) variables in the model, from the Peptide descriptor type option menu. The choices are:

- zvalue—Use the three z-value variables (z1, z2, z3) of Hellberg et al. [7] for the amino acid descriptors. These are derived from a principal components analysis (PCA) of 29 physicochemical variables for the 20 coded amino acids. The descriptors include molecular weight, pK_a, pI, side-chain vdW volumes, NMR shifts, retention times, partition coefficients, solvent exposure. Choose this variable set only if the peptides in your set consist entirely of coded amino acids.
- ezvalue—Use the five extended z-value variables of Sandberg et al. [8] for the amino acid descriptors. These are derived from a principal components analysis of 26 physicochemical descriptors for 87 amino acids (including the 20 coded amino acids). The descriptors include molecular weight, NMR shifts, partition coefficients, side-chain vdW volumes, HOMO and LUMO energies, heats of formation, polarizabilities, surface areas, hardnesses, TLC retention times, hydrogen-bond donor and acceptor counts, side chain charges.
- dpps—Use the 10 divided physicochemical property scores of Tian et al. [9]. These are derived from 23 electronic, 54 hydrophobic, 37 steric and 5 H-bond properties of the 20 coded amino acids, by applying principal components analysis to each of the groups separately and keeping 4 electronic components and 2 each for the other groups. Choose this variable set only if the peptides in your set consist entirely of coded amino acids.
- all—Use all of the above variables. There is likely to be some linear dependence between these variable sets.

The models are built using partial least squares techniques, with a choice of two variants that you can select from the QSAR method option menu. For each of these variants, you can set options to control the application of the method, by clicking the Advanced Options for *variant*, and making settings in the dialog box that opens—see Section 10.2.5 on page 64.

When you have finished making settings, click Build. The model may take a minute to build, and then the results are displayed in the Results tab. To apply this model to other sequences, you must save it first. Click Export Model in the Results tab to save the model to a file.

10.2.3 Examining the Model

When a model is built or applied, the predicted values and statistics for the prediction are shown in the Results tab. Model building actually builds multiple models for increasing numbers of partial least squares factors. To examine the results for a particular number of factors, set the number of factors using the PLS factors or KPLS factors box. The predictions and statistics are then shown in the tables below.

The Statistics tables display statistics for the training set and the test set. For the training set, four statistics are shown: the standard deviation, the R^2 correlation coefficient, the R^2 - correlation coefficient from cross-validation, and the stability. The latter two statistics are calculated with a leave-n-out method; the stability indicates how sensitive the results are to the choice of the training set. For the test set, the root-mean-square error, the Q^2 correlation coefficient and the Pearson r value are shown.

The results table shows the observed and predicted values for the peptide sequences used for the training and test sets, when building a model, and for the chosen sequences when applying a model. The table also shows the sequence name and the training and test set classification.

To use the results in another application, you can export them to a CSV file, by clicking Export Predictions and navigating to a location and naming the file in the file selector that opens.

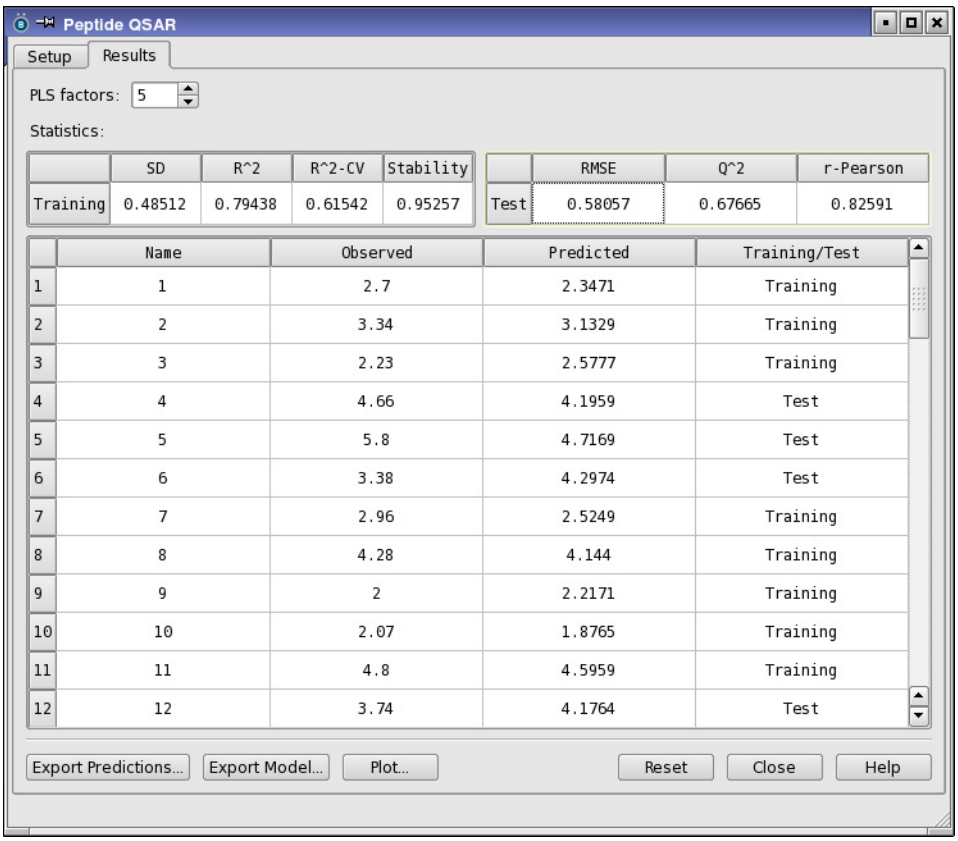

Figure 10.4. The Results tab of the Peptide QSAR panel.

If you have just developed a model, you must export it before you can apply it to a new set of sequences in the Setup tab. Click Export Model, and navigate to a location and name the file in the file selector that opens.

For a visual representation of the results, click Plot to display a scatter plot of predicted values against observed values, with a line of perfect fit, in the Peptide QSAR Scatter Plot panel. The panel has a plot toolbar, which allows you to configure the plot and save an image.

10.2.4 Applying the Model to Other Sequences

To apply an existing model, follow the steps below.

1. Load the sequences you want to apply the model to, if you do not already have the sequences loaded.

   See Section 10.2.1 on page 60 for information on loading sequences.

2. Choose the QSAR method for the model you want to apply.
3. Select Apply a model.
4. Click Browse to navigate to the model file and open it.
5. Choose an option for the table rows to apply the model to.

   You can apply it to all rows, or to rows that are marked in the table as training, test, or either, or to rows that are not marked.

6. Click Apply.

The model is applied and the results are displayed in the Results tab.

10.2.5 Setting Options for PLS Methods

The default options for the PLS methods usually give good results. However, if you want more control over the application of these methods when building a model, you can click Advanced Options for *method* and make settings in the dialog box that opens.

10.2.5.1 PLS Options

Set values for the standard PLS method with the following options:

Maximum number of PLS factors

Specify the maximum number of PLS factors to use in the regression model. Regression models are built for increasing numbers of PLS factors up to this number. The maximum number that can be used is limited by the number of descriptors, which is 3 times the number

of residues for the zvalue set, 5 times the number of residues for the ezvalue set, and x times the number of residues for the dpps set. It is rarely useful to build models with more than a few PLS factors, as models with a large number tend to be overfit.

Autoscale X variables

Scale the X variables by dividing the values of each property by the standard deviation in the value of the property.

Stop adding PLS factors when standard deviation of the regression drops to

Select this option to stop adding PLS factors when the standard deviation of the regression drops below the value specified in the text box. Using this option could result in fewer PLS factors than the number specified in the Maximum number of PLS factors box, but adding more factors may not yield any improvement in the model.

Eliminate X variables with t-value <

Eliminate X variables whose t-value is less than the value given in the text box. The t-value is the ratio of the coefficient of the variable in the fitted model to the standard error of the model. Small t values indicate that the variable is not contributing significantly to the model.

10.2.5.2 KPLS Options

Set values for the kernel-based PLS method with the following options. For more information on this method, see Section 3.5.3 of the *Canvas User Manual*.

Maximum number of KPLS factors

Specify the maximum number of KPLS factors to use in the regression model. Regression models are built for increasing numbers of KPLS factors up to this number. The maximum number that can be used is limited by the number of descriptors, which is 3 times the number of residues for the zvalue set, 5 times the number of residues for the ezvalue set, and x times the number of residues for the dpps set. It is rarely useful to build models with more than a few PLS factors, as models with a large number tend to be overfit.

Kernel nonlinearity

Change the kernel nonlinearity value. A Gaussian kernel $\exp(-d^2/\sigma^2)$ is used, where d is the Euclidean distance between two X variables. The nonlinearity value is $1/\sigma$, so small values are almost linear, and large values are very nonlinear. Higher nonlinearity typically leads to tighter fitting, but it also tends to give poorer predictions on new peptides. Click Reset to reset the value to the default.

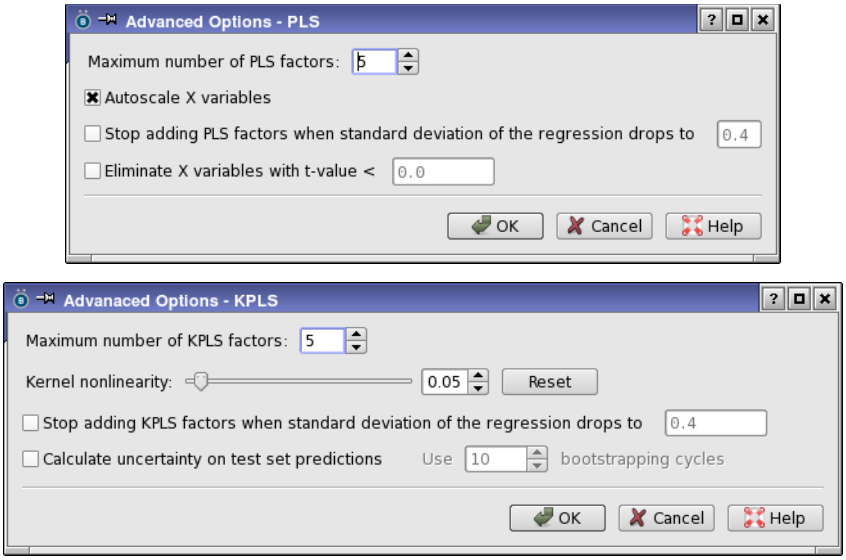

Figure 10.5. The Advanced Options dialog boxes

Stop adding KPLS factors when standard deviation of the regression drops to

Select this option to stop adding KPLS factors when the standard deviation of the regression drops below the value specified in the text box. Using this option could result in fewer KPLS factors than the number specified in the Maximum number of KPLS factors box.

Calculate uncertainty on test set predictions

Calculate a confidence interval for each predicted value in the test set, by bootstrapping. This is done by sampling the training set randomly with replacement to generate a new test set of the same size with duplicates, building a model and making predictions of the test set, then repeating the procedure a specified number of times. The standard deviation from the original test set is then calculated as the uncertainty.

Use N bootstrapping cycles

Specify the number of times a random sample is made and a prediction obtained in the uncertainty calculations. This number determines how many values are used in calculating the standard deviation, and should be at least 5.

10.3 Peptide Binding to a Receptor

As small peptides are of interest as therapeutic agents, it is useful to dock them to a receptor and examine the docked poses and a measure of their binding affinity. This can be done with the Peptide Docking panel, which you open by choosing Tasks → Peptide Docking.

Peptide docking is done with Glide, which treats the receptor as a rigid structure but with softening of the potentials in the active site region to simulate small adjustments of the receptor to the ligand. The peptide ligands are treated flexibly. As peptides are very flexible compared to typical non-peptide ligands, the Glide docking is performed with increased sampling, and several docking runs are done for each peptide, with different input conformations, to further increase the sampling.

10.3.1 Defining the Protein Receptor Region

The first task is to define the binding site for the peptides on the protein. The grids that are generated for the docking are centered on the location that you choose for the binding site. Two options are available:

- Centroid of Workspace ligand—Center the binding site and thus the docking grids at the centroid of the ligand molecule that you pick in the Workspace. Use this option if your receptor has a ligand that occupies the binding site. To pick the ligand molecule, select Pick and click on a ligand atom in the Workspace. Information on the ligand is shown in the box once you have picked the ligand.
- Centroid of selected residues—Set the center of the binding site at the centroid of a set of residues that you select. This option is useful if you don't have a bound ligand. You should choose residues whose centroid is approximately where the centroid of a bound ligand should be.

  To select the residues:

  1. Click Select Residues.
  2. Click the Residues tab in the Atom Selection dialog box.
  3. Pick the residues in the Workspace.
  4. Click Add, then click OK.

The X, Y, and Z text boxes show the coordinates of the center of the binding site.

Once you have defined the center of the binding site, you next need to define its extent, as Glide needs to know how large the volume is in which the ligand can be docked. This is calculated from the maximum number of residues in any peptide that you want to dock, which you specify in the Maximum peptide ligand length box.

With the Glide technology, the absolute maximum length of a peptide that can be docked is 16 residues. Glide also restricts the number of rotatable bonds to 50, due to the rapid increase in the number of conformations with the number of rotatable bonds. This limit means that the maximum peptide length that can be docked in practice is probably a lot less than 16, depending on the residue types.

10.3.2 Specifying the Peptides to Dock

The next task is to specify the peptides to dock, which can be taken from the Project Table or a file, or can be entered by hand. The sequences are listed in the peptide list. You can add peptides from multiple sources, with the Add sequences from buttons:

- File—Add peptide sequences from a file. Opens a file selector in which you can locate and open the file. The file must be a plain text file (`.txt`) or a PDB file (`.pdb`). If the file is a plain text file, there must be one sequence per line in the file. If it is a PDB file, only the first sequence is used.
- Project Table—Use the peptide sequences from the selected entries in the Project Table. You should select the entries in the Project Table first, before clicking this button. Only the sequences are used: no structural information is kept, as structures are generated by a conformational search.

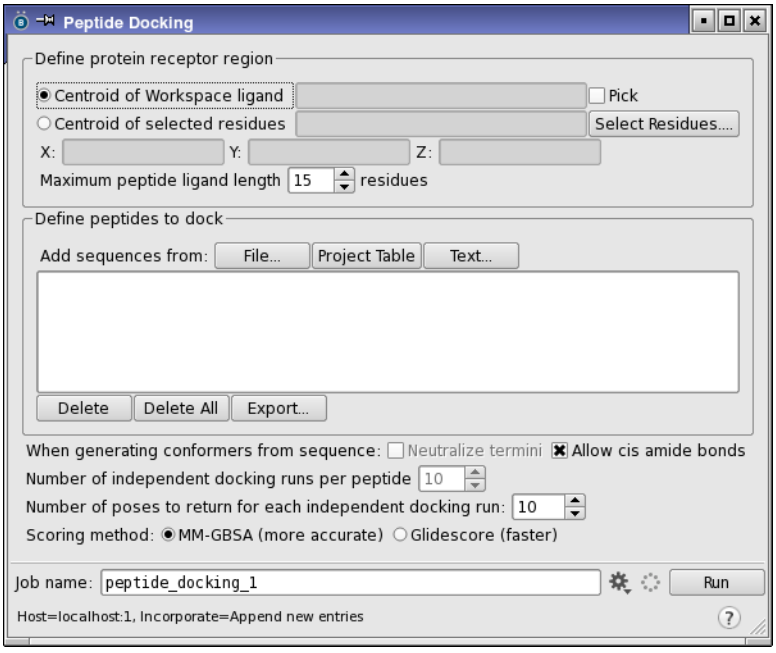

Figure 10.6. The Peptide Docking panel.

- Text—Enter the sequence of the peptide to be docked as a string of one-letter standard residue codes. Opens a dialog box in which you can type in the sequence.

The sequences are added to the list can be edited or removed. If the sequence is colored red, it contains an invalid single-letter residue code, and must be corrected. You can edit individual rows by double-clicking on the row, then editing the text. You can select multiple rows to remove them with the Delete button.

If you want to use the sequences somewhere else, you can export them to a plain text file, one sequence per line, by clicking Export.

10.3.3 Setting Docking Options

Most of the Glide docking options are set by the peptide docking protocol, or have their default values. A few relevant options are available in the panel.

The structures that are built do not have capping groups, and by default are in zwitterionic form, i.e. with charged termini. (The option to neutralize termini is turned off and cannot be changed in the current release.)

When generating conformers from sequences, you can allow peptide linkages in the cis conformation rather than the usual trans conformation by selecting Allow cis amide bonds.

Each peptide is docked several times in independent docking runs, with different starting conformations. This helps to increase the sampling of conformations in the docking process, and thus produce a better set of poses for the peptide. You can specify the number of independent docking runs to perform for each peptide in the Number of independent docking runs per peptide box.

After these runs have finished, the poses from each run are combined and sorted to give a set of poses for the peptide. You can specify the number of poses to return from each of these runs in the Number of poses to return for each independent docking run box. If the number specified is too small, you risk losing good poses. For example, if the first ten poses from a particular run are lower than any of the poses from the other runs, and you only allow one pose from each run, you would lose nine of the best poses.

Finally, you can set options for scoring the docked poses. The scores are estimates of the binding free energy, and are returned as properties of the docked poses. The two options are:

- MM-GBSA—Use the MM-GBSA ligand binding energy as the score. This option involves an extra calculation to evaluate the binding free energy in implicit solvent, and is therefore more expensive, but more accurate.
- Glidescore—Use the GlideScore value from the docking run to score the peptide poses.

This is the Glide SP docking score, and is produced automatically as part of the docking run.

For more information on the Glide docking process, see the *Glide User Manual*. For more information on the MM-GBSA method used, see Chapter 8 of the *Prime User Manual*.

10.3.4 Running the Job

When you have made all the settings, click Run to run the job with the current job settings, or click the Settings button to open the Job Settings dialog box and make settings for the host, distribution of the work, and return of the results.

Docking peptides is time consuming as there are many conformations to explore. A single peptide can take a few hours to dock. It it recommended that you run the job across multiple processors if you can, especially if you want to dock multiple peptides. The number of processors and the number of subjobs can be specified in the Job Settings dialog box.

The job output is a file containing the receptor and the poses of the docked peptides (in a "pose viewer" file). You can step through the poses in the presence of the receptor using the Pose Viewer panel—see Chapter 6 of the *Glide User Manual* for more information.

Protein-Protein Docking

The question of whether one protein binds to another, and where, can be addressed by protein-protein docking. Protein-protein docking in BioLuminate is performed using the Piper program [5], under license from Boston University. The job can be set up in the Protein-Protein Docking panel. In this panel you can set up jobs to dock two arbitrary proteins, dock an antigen to an antibody, or dock one protein to itself to form a dimer or a trimer.

One protein is treated as the "receptor" and the other as the "ligand". In the general case, it does not matter which protein is treated as the receptor and which protein is treated as the ligand. For antibody-antigen docking, the receptor is the antibody and the ligand is the antigen. The algorithm samples all possible orientations of the two proteins, subject to whatever constraints are applied. It uses a grid to locate the best poses of the two proteins, with a maximum resolution in the poses of about 5°. The docking is performed as a rigid-body optimization: there is no subsequent minimization of the interfacial region.

To open the Protein-Protein Docking panel, choose Tasks → Protein-Protein Docking.

For a tutorial introduction to protein-protein docking, see Chapter 10 of the *BioLuminate Quick Start Guide*.

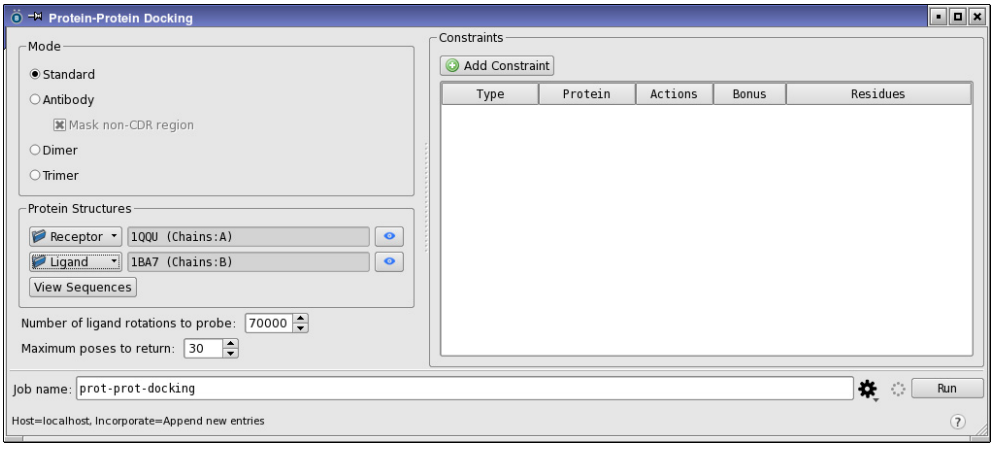

Figure 11.1. The Protein-Protein Docking panel.

11.1 Preparation of Proteins for Docking

Although it is not necessary to prepare proteins for docking, it may be advisable to do so. Proteins from the PDB often do not have coordinates for side chains on the surface of the protein, or even for whole chains that are solvent-exposed, as these are usually more mobile and it can be difficult to determine their coordinates from the X-ray data. If these missing parts of the structure are not in the favored binding region, the docking should produce good results. If the poorly defined parts of the structure are in the binding domain, pre-docking preparation of the structure can appreciably improve results. An example of the effect of missing side chains in the binding region is given in Chapter 10 of the *BioLuminate Quick Start Guide*.

To prepare your protein for docking, you can use the Protein Preparation Wizard (see the *Protein Preparation Guide*). When you prepare your structure, you should select Fill in missing side chains using Prime to predict the side chains, and Fill in missing loops using Prime to predict the loops. The loop prediction used in the protein preparation is the faster look-up method, rather than the more extensive ab initio loop building. Both of these predictions can take several minutes.

If you are concerned about the accuracy of the surface side chains or loops, you can do a more extensive prediction in the Refinement panel (Tasks → Loop and Sidechain Prediction). If you have access to the X-ray data, you could consider performing some refinement of the structure with PrimeX, the X-ray refinement program. Choose Tasks → Advanced Tasks → Protein X-Ray Refinement → Display Toolbar, and use the toolbar to perform the refinement tasks. See the *PrimeX User Manual* for more information.

In the docking experiment, however, the accuracy of the methods may not be sufficient to distinguish between conformations, so an extensive prediction of the surface side chains or loops is probably not necessary. The presence of the side chain or loop in the right region is likely to be more important than its exact conformation. If you want to test the effects of loop or side chain conformations, you can generate multiple conformations and dock each of them.

11.2 Docking a Protein to Another Protein

To dock a protein to another protein, without any special conditions on the type of protein, choose Standard in the Mode section of the Protein-Protein Docking panel.

Next, choose the receptor protein and the ligand protein. It does not matter which of your two proteins you choose as the receptor and which you choose as the ligand. To select the structures, click Receptor or Ligand in the Protein Structures section, and choose a source from the menu that is displayed. The menu items are the same for each menu:

- Browse for File—Opens a file selector so that you can browse to the location of the file and select it. The structure is added to the project and to the Workspace.
- From the Workspace—Use the structure that is in the Workspace. You should ensure that only the desired structure is included in the Workspace. If you have prepared the protein with the Protein Preparation Wizard, the structure should already be in the Workspace.
- From PDB ID—Opens a dialog box in which you can specify a PDB ID. The protein you specify is imported from the PDB, either from a local copy or from the web site. It is added to the project and to the Workspace. Proteins from the PDB are likely to have atoms missing.

When you import the structures, if hydrogens are missing, you are prompted to add them. If the protein has multiple chains, you are prompted to choose the chains. You can choose more than one chain for the docking.

When the protein is imported, the structure is added to the project and included in the Workspace. If the structure is removed from the Workspace you can add it with the Include and zoom button (eye icon).

To view the sequences of the proteins you imported in a sequence viewer, click View Sequences.

When you have selected the proteins to dock, you can set two parameters to control the docking. The first is the number of ligand orientations to the receptor that are sampled. You can set the value in the Number of ligand rotations to probe box The default of 70,000 corresponds approximately to sampling every 5° in the space of Euler angles, and is the maximum value allowed. Decreasing the number of rotations generally degrades the results, but decreases the run time.

The second parameter is the number of poses to return, which you can set in the Maximum poses to return box. Each pose is the center of a cluster that results from clustering the top 1000 results of rigid docking of the ligand. If more clusters are found than the number of poses to return, the clusters are ranked by size and poses are chosen from the largest clusters. If fewer clusters are found than the maximum number of poses to return, one pose per cluster is returned.

After setting the parameters, click Generate Models to run the docking job. A Start dialog box opens, in which you can choose whether to append the results to the project (Append new entries) or leave them on disk (Do not incorporate), choose a host to run the job on, name the job, and start it. Docking jobs are run on a single processor. A typical docking job with the default parameters takes several hours.

11.3 Applying Constraints

You can add constraints in the docking process, by providing an additional attractive term to the potential, remove the attractive potential, or declare residues to be buried, for residues that you select. Constraints are implemented as a bias during the docking process. They increase or decrease the likelihood of the specified interaction, but do not guarantee that the specified restraint will be met by all top ranking poses.

Constraints can be added in the Constraint section of the panel. Each constraint is represented by a row in the table in this section. To add a new constraint row to the table, click Add Constraint. All of the settings needed to define the constraint can be made in the cells of the table.

Table 11.1. Columns in the Constraints table.

Column	Description
Type	Choose the type of constraint to apply. The types are: Attraction—Increase the attractive potential for participation in binding. The value in the Bonus column is added to the default value of 1 to define the scaling factor for the attractive potential. Buried—Increase the van der Waals radius and the repulsive potential for buried residues, and set the attractive potential for ligand or antigen residues to zero. Repulsion—Set the attractive potential to zero, leaving only the repulsive van der Waals potential, which is not modified.
Protein	Choose the protein that the constraint applies to. The choices on the menu are the same as the labels on the button menus in the Protein structures section.
Actions	This column has two action buttons, one for deleting the constraint, and the other for zooming in on the atoms that define the constraint.
Bonus	Define the increase in the attractive potential for Attraction constraints. The value must be in the range 0.11 to 0.99. This value is added to 1 to define a scaling factor for the attractive potential.
Residues	Lists the residues affected by the constraint, and provides tools to select the residues. If you click in a table cell in this column, you can choose from two sources of the selected residues: From Workspace selection—Use the current Workspace selection to define the residues for the constraint. Choose Workspace atoms—Open the Atom Selection dialog box, so you can select the residues for the constraint. The selection you make is filled out to complete any residues that are only partly selected.

To define the constraint, you must choose the constraint type in the Type column, choose the protein to apply it to in the Protein column, set a value in the Bonus column if it is an attractive constraint, and then select the residues to apply the constraint to by using the Residues column. You can select the residues in the Workspace, then choose From Workspace selection, or you can choose Choose Workspace atoms, then use the Atom Selection dialog box to select the residues for the constraint. See Section 6.5 of the *Maestro User Manual* for information on using the Atom Selection dialog box.

To delete a constraint, click the red minus icon in the Actions column.

11.4 Creating Dimers and Trimers

You can dock a protein to itself to create a dimer or a trimer. To do so, choose Dimer or Trimer in the Mode section of the Protein-Protein Docking panel. When you do, only one of the menus in the Protein structures section is available, and it is labeled Monomer. You can select the monomer protein in the same way as for the general case, and set up and run the job in the same way, as described in Section 11.2.

The dimer and the trimer are subject to symmetry constraints: the dimer must have a twofold axis of rotation (C_2 axis), and the trimer must have a threefold axis of rotation (C_3 axis).

11.5 Docking an Antigen to an Antibody

The docking of an antigen to an antibody can be performed, with constraints for the target region of the antibody. To dock an antigen to an antibody, choose Antibody in the Mode section. If you want to prevent docking to the non-CDR region, ensure that Mask non-CDR region is selected. The receptor is the antibody and the ligand is the antigen.

As for the other types of docking, you can choose the antibody and antigen in the Protein structures section. After prompting you to add missing hydrogens, the antibody is analyzed to locate the CDR regions and determine which are the light and heavy chains. A progress bar is displayed while the analysis is done. When the analysis finishes, another dialog box opens, prompting you to select the chains to use for the antibody (receptor). You must select two chains: a light chain and a heavy chain. When selecting the antigen, an alert box opens, warning you about the chains. You can dismiss this box.

Apart from these changes, docking an antigen to an antibody is set up and run in the same way as a general protein-protein docking job.

Homology Modeling of Proteins

If you have a protein sequence and want to build a homology model of the protein, there are three ways that you can proceed in the BioLuminate interface.

Proteins for which you expect a high homology with the template and that require only a straightforward alignment to the template can be modeled in the Homology Model panel. In the default mode, any missing loops are predicted using a curated database of known loops in the PDB. This approach is very fast and a full homology model can typically be generated using this panel in 2-5 minutes. To open this panel, choose Tasks → Homology Modeling → Simple Homology Modeling. The use of this panel is described below.

Proteins where the homology is not as high or where alignment of the template and the query is required can be modeled either in the Multiple Sequence Viewer panel (opened with Tools → Multiple Sequence Viewer), or with the Structure Prediction panel (opened with Tasks → Homology Modeling → Advanced Homology Modeling). For information on using these panels, see the *Multiple Sequence Viewer* document for the Multiple Sequence Viewer and the *Prime User Manual* for the Structure Prediction panel.

If you are interested in homology modeling of antibodies, see Chapter 17.

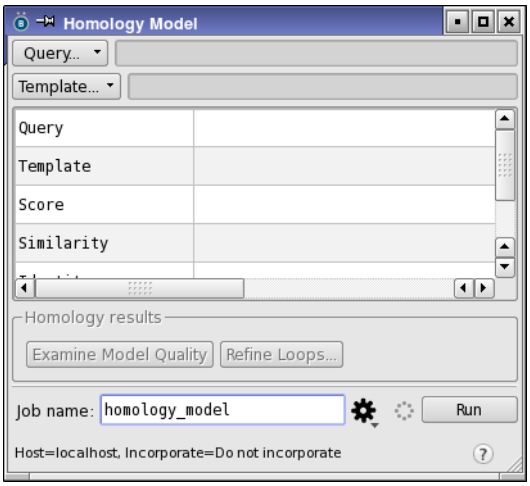

Figure 12.1. The Homology Model panel, initial view.

For a tutorial introduction to homology modeling using the Homology Modeling panel, see Chapter 6 of the *BioLuminate Quick Start Guide*. For a tutorial introduction to advanced homology modeling, see the *Prime Quick Start Guide*.

To build a homology model:

1. Choose a source for the query (or reference) sequence using the Query button. There are two choices:

   - From Workspace—Use the sequence of the structure in the Workspace as the query.
   - Browse for File—Opens a file selector so you can locate and import the sequence file for the query.

   When you have chosen a query, the box to the right of the button displays the text (Query structure) and the title of the query, if it has a title, or Query, if no title is available.

2. Choose a source for the template structure on which to build the model, using the Template button.

   You can either read in a template structure from a Maestro file or a PDB file (Browse for File), or you can run a BLAST search, as follows:

   a. Choose Template → BLAST Homology Search.

   b. Change settings if desired in the BLAST Search Settings dialog box, and click Start Job.

      When the job finishes, the Job Progress dialog box is replaced by the BLAST Search Results dialog box. This dialog box lists the homologs found with various measures of the match: E-value, Score, Identity, Similarity, and Homology.

   c. Choose a template from the list of homologs, and click Select Template.

      Blast search results are listed in order of decreasing homology score. By default, the first structure in the list is chosen, and in most cases this will be the most appropriate selection.

      If you do not have a local installation of the BLAST or PDB databases, a warning is displayed by default: “Multiple Sequence Viewer is attempting to access a remote server. Would you like to continue?” If you have not already turned this warning off, Select Do not ask this question again, to prevent it from being displayed each time a structure is downloaded, and click OK.

      The template is selected and the BLAST Search Results dialog box closes.

   When you have chosen a template, the box to the right of the button displays the text (Template structure) and the title of the template. The table rows in the Homology Model panel are filled in with information on the template.

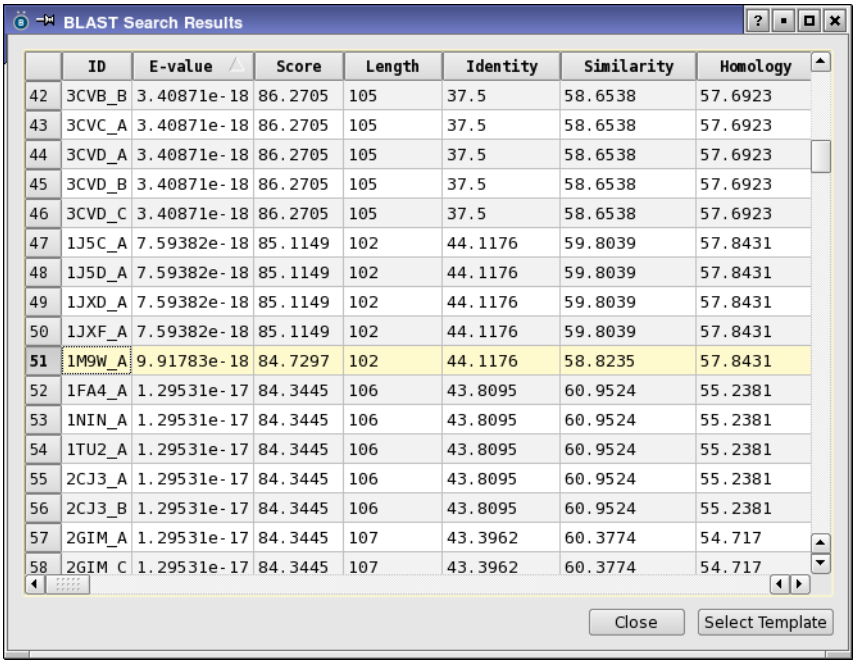

Figure 12.3. The BLAST Search Results dialog box.

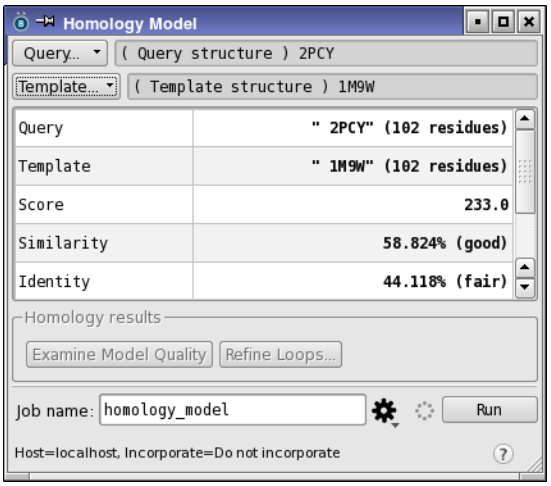

Figure 12.2. The Homology Model panel after selecting a template.

3. Click Generate Models.

   The Homology Model - Start dialog box opens. You can name the job and select a host to run it on. The job includes alignment of the template and the query using ClustalW, and the structure is built on the basis of the template and an analysis of structural elements in the PDB for non-templated regions (a "knowledge-based" selection of the coordinates).

When the job finishes, the model is added to the Workspace, in cartoon representation. The cartoon is colored by how the template was used: dark blue for residues for which all coordinates were taken from the template, cyan for residues for which the backbone was taken from the template, and red for residues that were entirely modeled, not using the template. The title given to the model includes information on the query and the template.

If the job fails, it is likely that there is insufficient homology or poor alignment between the reference and the template to build a model. In this case you should use the Advanced Homology Modeling panel (Tasks → Homology Modeling → Advanced Homology Modeling) or the Multiple Sequence Viewer (Tools → Multiple Sequence Viewer).

If you want to examine the quality of the model structure, click Examine Model Quality, to open the Protein Structure Quality Viewer panel, where you can view reports on the protein structure, a Ramachandran plot, and plots of protein properties

If you want to refine loops in the model, click Refine Loops. The Refinement panel opens with the Refine loops task selected. You should only need to refine the loops if they were not predicted from the template. Click Non-Template in the Refinement panel to list only the loops that did not come from the template. See Chapter 6 of the *Prime User Manual* for information on this panel.

Residue and Loop Mutation

At some point in a workflow, you might want to mutate a single residue, or replace a single loop with another loop. You can do this in the Residue and Loop Mutation panel, which you open from the Tasks menu.

BioLuminate provides other tools for mutations of more than one residue. The Residue Scanning panel allows you to mutate a protein at multiple sites to generate a set of proteins, each with a single mutation—see Chapter 15 for information. If you want to mutate residues to or from cysteine and break or form disulfide bonds, you should use the Cysteine Mutation panel—see Chapter 16 for information.

The workflows in this panel are described in the following sections.

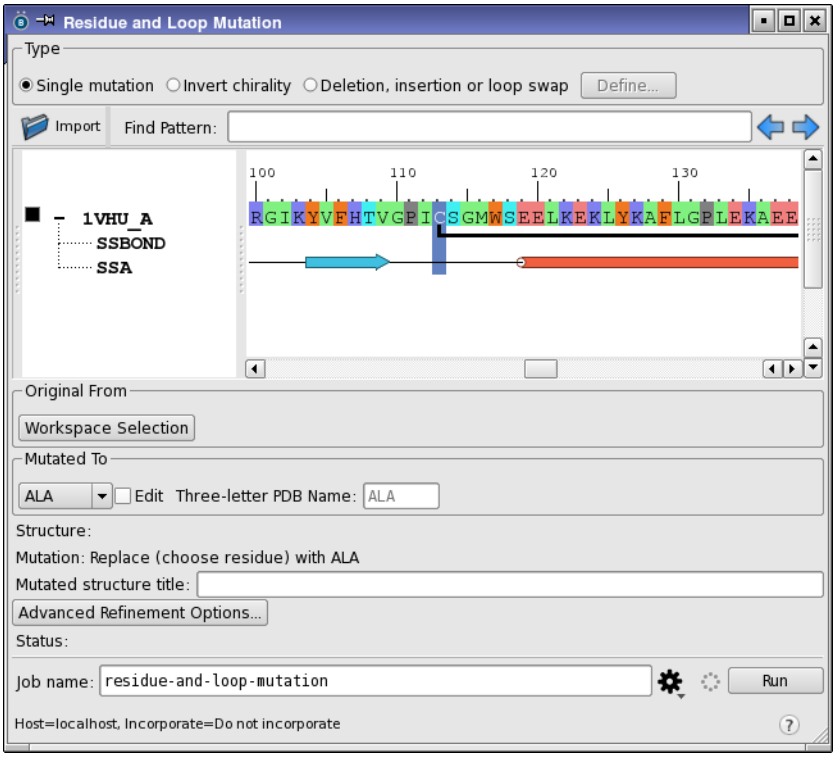

Figure 13.1. The Residue and Loop Mutation panel.

13.1 Residue Mutations

There are three types of single-point mutation that you can make in this panel: mutation to a standard amino acid, mutation to a nonstandard amino acid, and chirality inversion. You can choose which type to do with the options in the Type section. Select Single mutation for both standard and nonstandard amino acids; select Invert chirality to invert the chirality of the amino acid.

The workflows for each of these mutations is summarized here and detailed in the following sections.

To mutate a single residue to a standard amino acid:

1. Select Single mutation in the Type section.
2. Select the residue to be altered either in the Workspace or in the panel's sequence viewer.
3. Click Workspace Selection in the Original from section.
4. Choose the amino acid from the option menu in the Mutated to section.
5. Click Mutate.

To mutate a single residue to a nonstandard amino acid:

1. Select Single mutation in the Type section.
2. Select the residue to be altered either in the Workspace or in the panel's sequence viewer.
3. Click Workspace Selection in the Original from section.
4. Choose a standard amino acid from the option menu in the Mutated to section.
5. Select Edit.

   The Workspace structure is replaced by the chosen standard amino acid with the side-chain atoms shown in line (wire frame) representation.
6. Edit the side-chain atoms in the Workspace to produce the desired nonstandard amino acid.
7. Deselect Edit.
8. Provide a three-letter PDB residue name in the dialog box that opens.
9. Click Mutate.

To invert the chirality of a single residue:

1. Select Invert chirality in the Type section.
2. Select the residue to be altered either in the Workspace or in the panel's sequence viewer.
3. Click Workspace Selection in the Original from section.
4. Click Mutate.

13.1.1 Selecting the Residue to Mutate

You can select the residue to mutate by picking it in the Workspace structure or in the sequence viewer, then clicking Workspace Selection in the Original From section to register your choice.

To find a particular residue in the Workspace, you can use the Find tool. Type CTRL+F (⌘F) in the Workspace to display the Find toolbar below the Workspace. You can then choose Residue number or Residue type from the Find menu to choose residues by number or type, then enter the number in the text box or choose the type from the menu, and click the N or P button.

To find a residue in the sequence viewer in the Residue and Loop Mutation panel, enter the residue letter in the Find Pattern text box, and click the arrow keys to step through the occurrences. You can also enter multiple residues to find a pattern, e.g. `D-A-P`. The pattern uses PROSITE syntax, which is explained in the tool tip. When you have found the residue you want, clear the Find Pattern text box, then click on the residue to select it, or simply select it in the Workspace. The residues that are found are selected in the Workspace, so clicking on one of them changes the selection to the desired residue.

If the sequence viewer doesn't have a sequence in it, you can click Import and choose From Workspace to load the sequence of the Workspace structure.

13.1.2 Defining the Mutation

When you have selected the residue to mutate, you can define the mutated residue. The controls to do so are only available if you selected Single mutation in the Type section. If you selected Invert chirality, the mutation is already defined. You can mutate to a standard amino acid or a custom amino acid. In both cases, you start by choosing a standard amino acid from the option menu.

To begin editing a standard amino acid in the Workspace to produce a custom amino acid, check Edit. The amino acid replaces the protein in the Workspace, with the backbone represented as ball and stick and the side chain as lines.

To add groups to the structure or to change elements, you can use the Build and the Fragments toolbars. Click Build or Fragments on the Manager toolbar at the top of the main window to

display these toolbars. With the Fragments toolbar, you can select a fragment, then click on an atom to replace that atom with the fragment. With the Build toolbar, you can sketch a structure with the Draw tool, change the element with the Set Element tool. You might also want to add hydrogens after sketching a structure, which you can do from the Edit toolbar with the Add H tool. You should also use the Clean Up tool after sketching the structure and adding hydrogens, to ensure that the structure is not distorted.

For more information on building structures, see Chapter 5 of the *Maestro User Manual*.

When you have finished editing, clear the Edit check box. A dialog box prompts you to provide a 3-letter PDB name for the custom amino acid. The default is USR. You can edit the name in the Three-letter PDB Name text box after creating a custom amino acid.

The text on the amino acid option menu is set to Custom when you finish editing.

After the mutation is defined, the mutation is reported in the panel, and a title for the mutated structure is entered in the Mutated structure title text box. You can edit this title if you wish.

13.1.3 Refining the Mutated Structure

If you want to refine the mutated structure, click Advanced Refinement Options and make selections in the Refinement Options dialog box. It is usually a good idea to do some sort of refinement, to allow the protein to adjust to the new residue. This is particularly so if you created a custom residue, which might not be in the most favorable conformation.

There are two types of refinement available: minimization and molecular dynamics simulation. You can choose one or the other or both. The minimization is run first.

Minimization is the fastest option, and is a good choice if the mutated residue is in approximately the right conformation. If you choose to do a minimization, you can run it in the gas

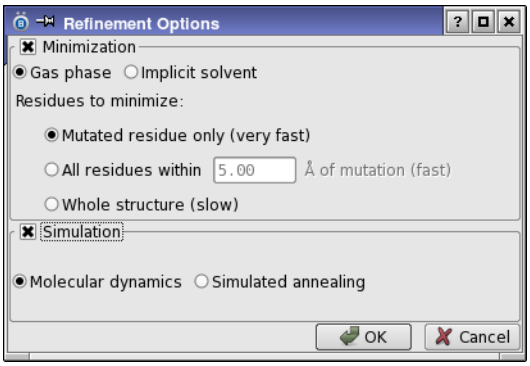

Figure 13.2. The Refinement Options dialog box.

phase or in implicit solvent. The residues minimized can be limited to just the altered residues, all residues within a given distance of the altered residues, or the entire structure.

Molecular dynamics is much more time-consuming, but it can sample other parts of conformational space, particularly if you run simulated annealing. Once the structure is mutated and any minimization is done, the protein is prepared for the simulation by adding explicit water molecules. The Molecular Dynamics panel or the Simulated Annealing panel opens, and you can make settings and run the simulation, which uses the Desmond molecular dynamics program. For more information on these panels, see Chapter 3 of the *Desmond User Manual*.

If you decide you do not want to run the simulation, which can take many hours, you should delete the entry group in the Project Table that was created for the simulation, as follows:

1. Select the entry group (click the row with the number in square brackets).
2. Choose Entry → Unlock.
3. Choose Entry → Delete.

For information on using the Project Table, see Chapter 9 of the *Maestro User Manual*.

13.2 Insertions, Deletions, and Loop Swaps

You may want to perform larger structural changes than a single residue mutation, such as deleting or inserting multiple residues, or replacing a loop with another loop. To do this, select Deletion, insertion, or loop swap in the Type section, then click Define to define the changes to be made in the Insertions, Deletions, and Loop Swaps panel.

The basic procedure is summarized below, and details are given in the following sections.

To insert or delete residues or modify a loop:

1. Select Deletion, insertion or loop swap in the Type section.
2. Click Define.

   The Insertions, Deletions, and Loop Swaps panel opens.
3. Choose the loop to be modified.

   For best results select at least two residues on either side of the residues to be modified. The selection can be made either by using the sequence viewer or by selecting residues in the Workspace.
4. Click the Workspace Selection button to load the original loop into the table.

5. If desired, load a new structure and select residues to specify a starting set of residues for the replacement loop.

   By default, the replacement loop starts out identical to the original loop.

6. If desired, edit the replacement loop by clicking on residues in the table and specifying an insertion, deletion or mutation at that point.

7. Choose the number of models to generate and the prediction method.

8. Click Accept.

   The Insertions, Deletions, and Loop Swaps panel closes.

9. Click Mutate.

No refinement of the loop modification is offered, so you must perform any refinement independently.

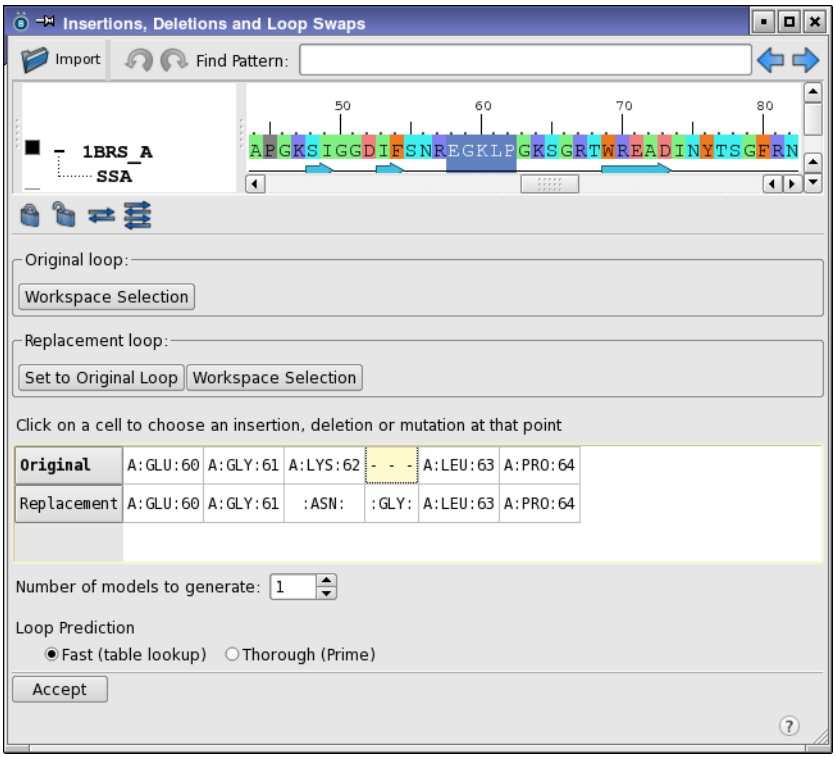

Figure 13.3. The Insertions, Deletions, and Loop Swaps panel.

13.2.1 Choosing the Original and Replacement Loop

When the Insertions, Deletions, and Loop Swaps panel opens, the sequence viewer at the top is loaded with the sequence of the structure in the Workspace. If you want to use a different s structure, you can replace the Workspace structure with another structure, or you can import a structure from a file or from the PDB. To load another structure, click Import, and choose the source from the menu that is displayed. The choices are:

- Browse for File—Open a file browser in which you can navigate to the desired location and select the file that contains the structure. The allowed file types are Maestro and PDB.
- From PDB ID—Import the structure from the specified PDB ID. Opens the Enter PDB ID dialog box, in which you can enter the PDB ID of the sequence. The sequence and structures are retrieved from a local copy of the PDB if it is available, or from the RCSB web site, depending on the preference set for PDB retrieval.
- From Workspace—Import the structure that is displayed in the Workspace.

The sequence for the imported structure is added to the sequence viewer. This allows you to import multiple structures, and choose which structure to take a loop from.

When you have the structure that you want to mutate, you can select the loop to mutate either in the Workspace or in the sequence viewer in the panel.

Selecting the loop in the Workspace allows you to visually identify the loop. To ensure that you are picking residues in the Workspace to define the loop, choose Edit → Pick Mode → Residues, or type the letter R with the pointer in the Workspace. You can then pick residues that belong to the loop you are interested in. The residues that you pick are highlighted in the Workspace (yellow dots) and in the sequence viewer (white letter on blue background).

Selecting the loop in the sequence viewer allows you to easily search for patterns in the sequence, or to use the secondary structure assignment (SSA) to identify loops. The SSA has no annotation where there are loops (the absence of any other secondary structure). To search for patterns in the sequence, enter the pattern in the Find Pattern text box, with a dash separating each residue from the next, e.g. `D-A-P`. The syntax is summarized in the tool tip for the text box, and is given in more detail on page 40. You can use the arrow keys to step through the patterns. All of the matches are selected in the Workspace, so you might have to select the loop that you want in the Workspace, or by clearing the Find Pattern text box, then selecting the residues to use.

When you select the residues for the loop to mutate, you should select at least two residues on either side of the residues that you plan to modify. These extra residues are “stem” residues, which are used by the Prime software when rebuilding the loop, to properly fit the new loop

onto the structure. For example, if a single residue is being deleted, the original loop should consist of five residues - the residue to delete and the two residues on either side of it.

After you have selected the residues, click Workspace Selection in the Original loop section, to register the selection and copy the structure for the loop mutation job. The original loop is used by default for the replacement loop, which you can modify as described in the next section.

If you want to replace the loop with a loop from another structure, you can place another structure in the Workspace, and select a loop from this structure. If the structure is not already in the sequence viewer, you can import it as described above. You can also select the loop in the same way as for the original loop.

You might want to align the sequence for the replacement to the sequence for the original structure, so that you can select the loop by its alignment. To do the alignment, make sure that the original sequence is at the top of the sequence viewer (use the shortcut menu for the sequence name to move it if necessary), and add the replacement sequence to the selected sequences. You can then use either the pairwise alignment tool or the multiple alignment tool to align the sequences (using ClustalW).

You can also do manual alignment, with locking and unlocking of gaps. See the *Multiple Sequence Viewer* document for more information on doing manual alignment.

When you have selected the residues in the desired structure for the replacement loop, click Workspace Selection in the Replacement loop selection. If you decide not to use this loop, you can click Set to Original Loop to change the replacement loop back to the original loop structure.

13.2.2 Editing the Replacement Loop

If you want to modify the replacement loop residues to create the new loop that will replace the original loop, you can do so in the table in the lower part of the panel. Clicking on a table cell opens a menu that allows you to delete or mutate the current residue or insert a new residue before or after the current residue. The change is applied to the Replacement row, regardless of whether you define the change for the original row or the replacement row. The Insert Before, Insert After, and Mutate items display a list of 20 standard amino acid residues that you can select.

When you choose to insert a residue, a new column is added to the table before or after the current position. The cell in the Original row has three dashes, to indicate that there is no residue in the sequence at this position. The cell in the Replacement row has the residue name

between two colons, to indicate that the chain and residue number are not yet assigned. Insertion codes are added for the inserted residues when the job is run.

If you delete a residue, the cell in the Replacement row is set to three dashes, to indicate the deletion.

Likewise, if you mutate a residue, the cell in the Replacement row has the new residue name between two colons.

13.2.3 Choosing the Output and Method

The replacement loop is built with the Prime structure building software. By default, only the best structure is returned, but if you want to examine more than one structure, you can specify the number of new structures to be returned in the Number of models to generate box.

You can also choose whether to create the loops via a fast table-lookup method or via a full Prime loop prediction, which builds loops by sampling multiple conformations and scoring them, to produce the best loop structures. The fast method takes only a minute or so, where as the thorough loop sampling can take hours. To make the choice, select Fast or Thorough under Loop prediction. No additional refinement beyond these methods is done for new loop predictions.

When you have finished selecting options, click the Accept button to return to the main Residue and Loop Mutation panel. Click the Mutate button on that panel to begin the loop swap job.

Cross-Linking Proteins

As part of protein design, it can be useful to cross-link two proteins. For example, suppose you have two oligopeptide fragments or protein domains that bind to a third protein. Both fragments or domains need to bind to the third protein for function. To increase efficacy, you might want to try to tether those two fragments or domains together. Another use for cross-linking is for circular permutation of a protein, in which you connect the termini and break the chain at some other point. Provided the break is outside the binding region, you could create a protein that still binds a ligand, but may interact differently in the cellular environment.

The Crosslink Proteins panel allows you to cross-link two pre-positioned proteins by connecting chain termini with peptide linkers. To open this panel, choose Tasks → Crosslink Proteins.

14.1 Preparing the Proteins for Cross-Linking

The link between the proteins is formed with standard peptide bonds, so the links must be formed from the N-terminus of a chain to the C-terminus of a chain. No adjustment of the relative position or orientation of the proteins is done in the process, so you must ensure that they are properly positioned before linking them—for example, if two proteins bind to a third and you want to link those two proteins, you may want to take their positioning from the complex with the third protein.

Before you can add the linkers, you must ensure that the proteins are in a single project entry. If they are in different entries, you can create a project entry by clicking the Workspace button on the Manager toolbar to show the Workspace toolbar, then click the Create Entry button, and name the entry in the dialog box that opens.

As for any modeling exercise, you should ensure that the protein is prepared, by using the Protein Preparation Wizard. To cross-link the proteins, you must delete all het groups and waters when you prepare the protein, and ensure that the protein contains only on the standard amino acids. The het groups can be restored later, for example by creating another project entry that contains only the het groups and waters, then merging this entry with the results of the cross-linking.

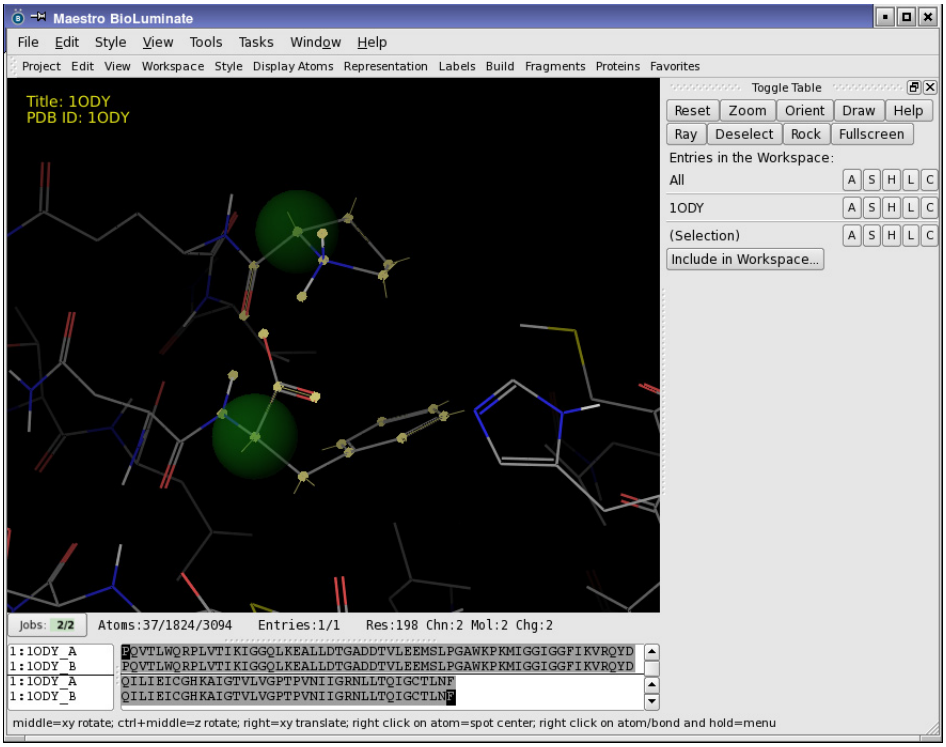

Figure 14.1. The Workspace showing selected termini.

14.2 Picking the Connection Residues

The first task is to pick the residues on the proteins that you want to connect with a linker. These residues must be protein termini.

When you pick the termini, it can be useful to select them in the Workspace first. To do this, you can display the sequence viewer (Edit → Settings → Show Sequence Viewer), select the terminal residues in the sequence viewer, then choose A → Zoom in the Selection row in the Toggle Table. The residues are marked with yellow selection markers, which makes them easy to pick.

To pick the first terminus, select Pick residue in Workspace for Connection residue one, and pick a terminal residue in the Workspace. A warning is posted if you pick a non-terminal residue. The text box is filled in with the residue ID in the form *chain*:*resnum*(*resname*), e.g. `A:1(PRO)`. The alpha carbon of the residue is marked with a green sphere in the Workspace. After the residue is picked the check box is automatically cleared.

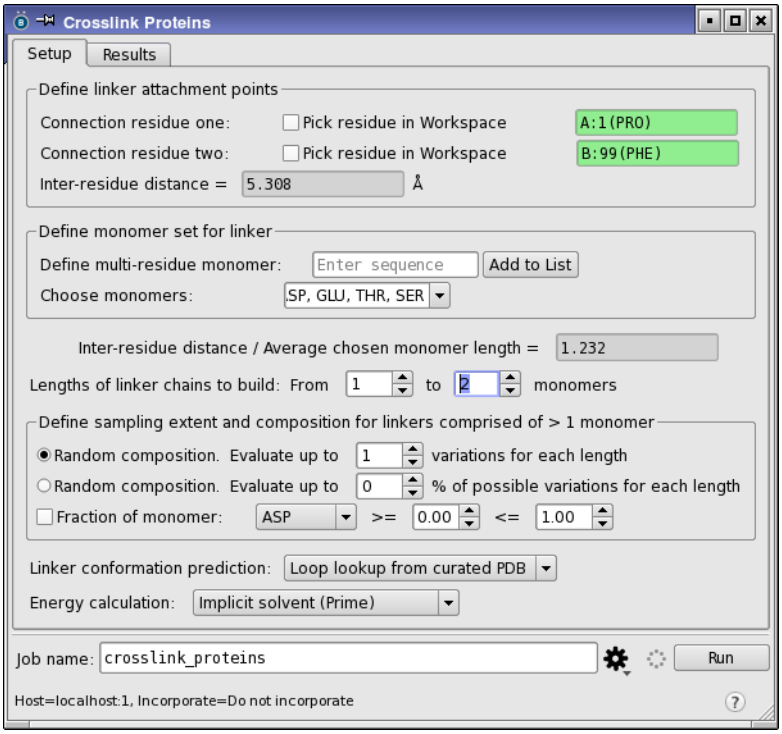

Figure 14.2. The Setup tab of the Crosslink Proteins panel.

To pick the second terminus, select Pick residue in Workspace for Connection residue two, and pick a terminal residue in the Workspace. If residue one was an N-terminal residue, residue two must be a C-terminal residue, and vice versa. A warning is posted if you pick a non-terminal residue or a terminus of the wrong type. The text box is filled in with the residue ID as for the first residue, and the alpha carbon of the residue is marked with a green sphere in the Workspace. When the residue is picked the check box is automatically cleared, and the two text boxes are colored green to indicate that the definition of the connection residues is complete.

After the two residues are picked, the Inter-residue distance text box is filled in with the distance between the alpha carbons of the two picked residues. The distance is used to calculate and display the approximate number of linker units needed to link the proteins.

14.3 Defining the Linkers

A great deal of flexibility is provided for the definition of the linkers used to connect the chain termini. The linkers are multimers, composed of monomers that are built from standard amino acids. By default, the standard amino acids are available as monomers, but you can define your own multi-residue monomers. You can also choose different numbers of monomer units to include in the linker chain. So, for example, if you want to specify the linkers completely, you can add monomers of the desired length, and allow only one monomer in the linker.

To define a monomer, enter the sequence of 1-letter codes for the monomer in the Define multi-residue monomer text box, and click Add to List.

The next step after defining the monomers is to select the monomers you want to use, from the Choose monomers menu. The monomers that you define are automatically selected, and their sequence is displayed in the text box of the menu. To add a monomer, choose it from the menu. The selected monomers have a check mark next to them. To remove a monomer, choose it from the menu; the check mark is removed and the sequence is removed from the text box. You can choose monomers of different lengths to create linkers with different numbers of residues.

Once you have chosen monomers for the linkers, the Inter-residue distance / Average chosen monomer length text box reports the ratio of the inter-residue distance to the average length of the monomers chosen for the linker. This ratio gives an approximate number of monomers that must be included in the linker to form a proper link. You can set the minimum and maximum number of monomers in the linkers on the basis of this information, in the Lengths of linker chains to build boxes.

The monomer units in each linker are selected randomly for each linker that has more than one monomer. The number of linkers to construct for each number of monomers can be specified explicitly or as a percentage of the possible variations, by choosing one of the Random composition options, and providing the number or percentage in the box. The total possible linkers scales as N^M, where N is the number of monomers chosen and M is the linker length.

The random selection can be modified by specifying the fraction of each monomer that you want in the linkers. Select Fraction of monomer, choose the monomer from the option menu, and set the limits in the text boxes. Linkers for which the fraction of each monomer does not fall within the specified range are discarded. No checking is done that the fractions add up to 1, so you must ensure that your choices do not result in impossible requirements (such as setting the minimum for two monomers greater than 0.5). A warning is presented if no linkers are generated, so if you see this warning, you should check the fractions, if you have set them.

14.4 Choosing Methods and Running the Job

The linkers are added by building a loop between the two selected termini. There are two methods available from the Linker conformation prediction option menu for determining the loop conformation:

- Loop lookup from curated PDB—Use a table of known loops taken from the PDB to determine the loop conformation. This is the fastest method.
- Simple de novo loop creation—Build the loop residue by residue to produce a single loop conformation that has no clashes with the existing structure.

It is a good idea to refine the structure after it is built, as the simpler methods used here might not give the optimal conformation, though a minimization of the linker is performed as part of the procedure.

The chains are connected with each linker that is generated, and the strain energy for the linker is evaluated as the difference between the linker energy in the linked conformation and the minimum energy of the linker in the unbound conformation. You can choose a method for calculating the strain energy of the new loop from the Energy calculation option menu.

To start the job, click Run, or click the Settings button and make settings in the Job Settings dialog box, then click Save and Start. The time taken depends on the length and number of linker chains.

14.5 Examining the Results

When the job finishes, the results are automatically listed in the Results tab. The results table lists the structures for each linker, showing details of the linker and its composition, the total strain energy, and the strain energy per amino acid. You can sort the table by the values in any column by clicking on the column heading.

The strain energy is useful to score multiple possible cross-linking chains on a relative basis. Generally, chains with lower strain energies are better. Performing a search of multiple lengths and conformations and focusing on those cross-links with the lower strain energies helps to select those candidates that are more likely to accommodate the connected domains in the starting conformation. However, note that the strain energy is only one component of the energy of the linked protein structure: it does not include the interaction energy between the linker and the protein, which can compensate for the strain.

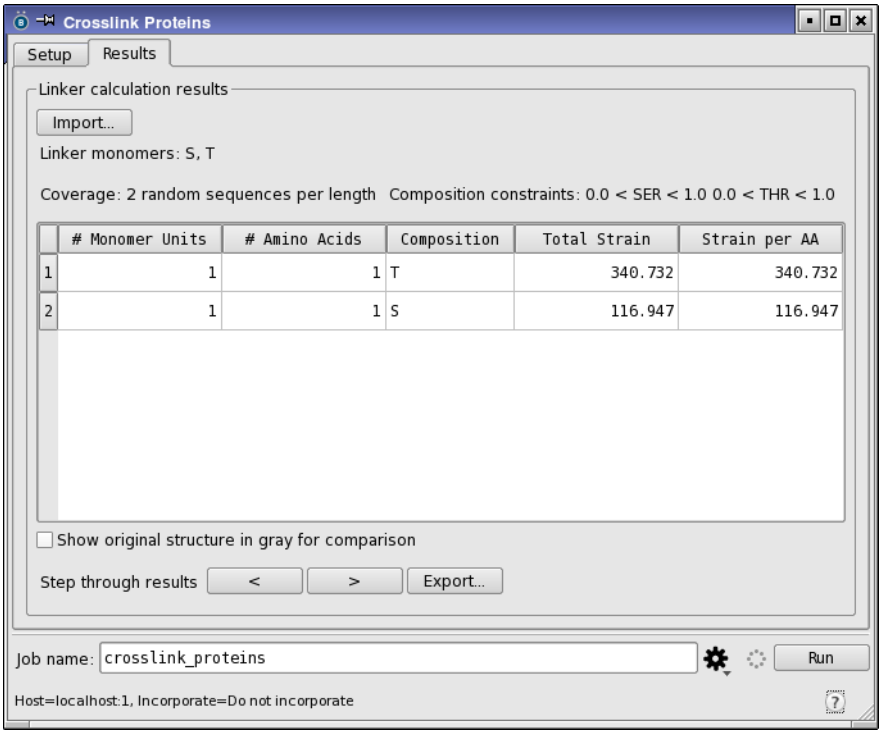

Figure 14.3. The Results tab of the Crosslink Proteins panel.

If you want to examine the structures, use the Step through results buttons, which display the structures in turn in the Workspace. The table row for the structure in the Workspace is highlighted. You can also add the original structure to the Workspace by selecting Show original structure in gray for comparison.

If you want to display the results at some other time, you can click Export, and export the results to a zip file. The results are also present in the Project Table. To view the results the Results tab later, you can click Import, and navigate to the zip file.

Scanning for Residue Mutations

This chapter describes how to scan a protein for potential residue mutations, generate mutated structures, and compare the properties of the mutated structures. The mutation sites and the mutations can be selected manually or selected automatically based on homology modeling or 3D structural criteria.

Some examples of the use of residue scanning are:

- improvement of protein-ligand affinity
- identifying protein-protein interface hotspots
- identifying residue mutations that can improve stability
- mutating unpaired, solvent-exposed Cys residues to reduce undesired reactivity

There are two panels available performing residue mutations. If you want to perform a set of single mutations in a single chain of the protein, you can use the Residue Scanning panel. If you want to perform several mutations in the protein and search for the optimum binding affinity or protein stability, you can use the Affinity Maturation panel. The panels have a very similar structure as the tasks to be performed are in most cases the same. The sections below describe each of the tasks, with notes about the differences between the two panels. The panels also use a common results viewer, which is described at the end of this chapter.

To open the Residue Scanning panel, choose Tasks → Residue Scanning → Perform Calculations in the main window.

To open the Affinity Maturation panel, choose Tasks → Affinity Maturation → Perform Calculations in the main window.

15.1 Selecting and Analyzing the Protein

The first step is to select the protein and to display it in the Workspace. The protein must be one that has been prepared for use in modeling. If it has not been prepared, we recommend that you prepare it with the Protein Preparation Wizard (on the Tools menu and the Tasks menu). Details of preparing a protein can be found in the *Protein Preparation Guide*.

If you have not already opened the panel, open it from the Tasks menu, as described above. You can display the protein before or after opening the panel.

The first part of the procedure is to analyze the protein, which you do by clicking Analyze Workspace. In this process, you may be asked some questions:

- If your structure has not been prepared for modeling, you are asked if you want to use the Protein Preparation Wizard to prepare the structure. You cannot proceed if you do not have a protein that is an all-atom, 3D structure with bonding information. After preparing the protein, display it in the Workspace again and click Analyze Workspace.
- If your structure has multiple chains, a dialog box opens.
  - For residue scanning, you are prompted to choose a chain. Mutations can only be performed on a single chain in a residue scanning run.
  - For affinity maturation, you are prompted to divide the chains into two groups. These groups are used to calculate the binding affinity, so one of them can be regarded as the "ligand" and the other as the "receptor". All chains can be mutated.

When the structure has been analyzed, the residues table is filled in with all the residues in the chain or chains that can be mutated.

Instead of analyzing an entire protein or an entire chain, you can select residues and analyze only the selected residues. To do this, select the residues in the Workspace structure, select Analyze only selected Workspace residues in the panel, then click Analyze Workspace.

15.2 Setting Up and Running the Job

The tasks involved in setting up the job are done from the Residues tab: selecting residues, selecting mutations, setting parameters for refinement. The Homology Suggestions tab provides a way of automatically selecting residues—see Section 15.3 on page 104—but the remaining tasks must be done in the Residues tab.

15.2.1 Choosing Residues to Mutate

The residues that are available for mutation are listed in the Residues for consideration table in the Residues tab. The residues are identified by the chain name, the residue number (and insertion code), and the 3-letter residue name. If you want to show only the polar or the nonpolar residues, select Polar or Neutral underneath the table. To redisplay all residues, select All.

The table also contains information on van der Waals surface complementarity[10] for residues that are at the interface between chains. This allows you to choose residues to mutate based on their complementarity.

The residues that are mutated when you run the mutation job are the residues that have mutations defined, by default. If you want to define mutations for multiple residues but don't want to mutate all of them in a given run, select the residues that you want to mutate in the table, and then select Mutate selected residues only. When the job is run, only the selected residues are mutated.

15.2.2 Choosing the Mutations

There are two main ways of choosing the mutations: manually, or based on analysis of homologs to the structure of interest. Manual selection is described in this section, and using homology is described in Section 15.3 on page 104.

You can choose mutations for individual residues, or you can apply a set of mutations to selected residues.

To define the mutations for a single residue:

1. Click in the Mutations column for the residue.

   An option menu is displayed in the table cell.

2. Choose the residue types from the option menu.

   The residue types include groups such as Neutral, Polar, Hydrophobic, Positive, and Negative. You can choose more than one residue or group from the list, and the new residues

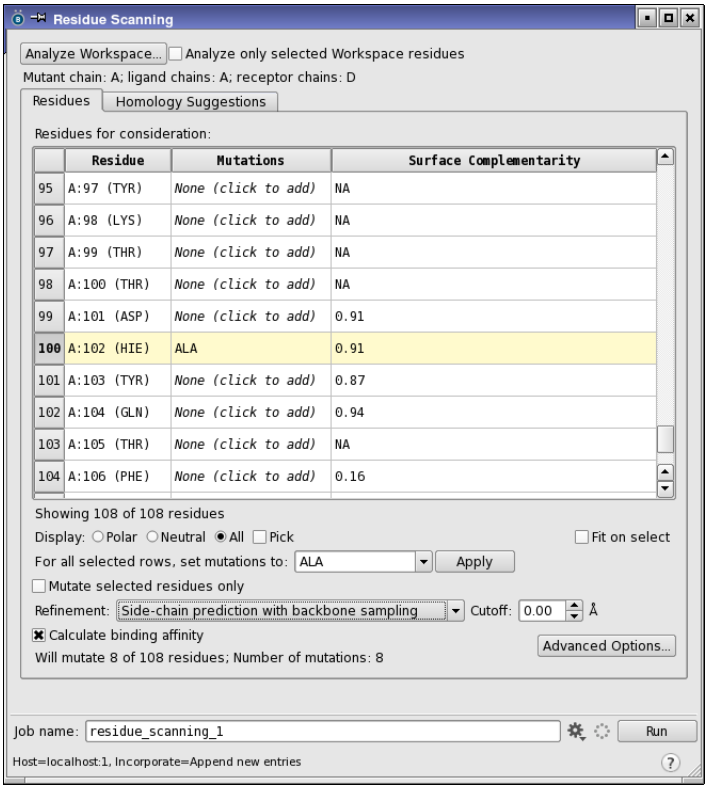

Figure 15.1. The Residue Scanning panel, Residues tab.

are added to the list. Each residue type that you select is checked on the list. The mutations are displayed in the box at the top of the menu.

3. Press ENTER to finish adding mutations and dismiss the option menu.

   The mutations are now listed in the table cell.

   If you want to cancel the changes and dismiss the option menu, press ESC.

To apply a set of mutations to selected residues:

1. Select the residues in the table that you want to apply the same mutations to.

   You can select Pick and pick residues in the Workspace to select them in the table.

2. Choose the residue types that you want to use for the mutations from the For all selected rows, set mutations to menu.

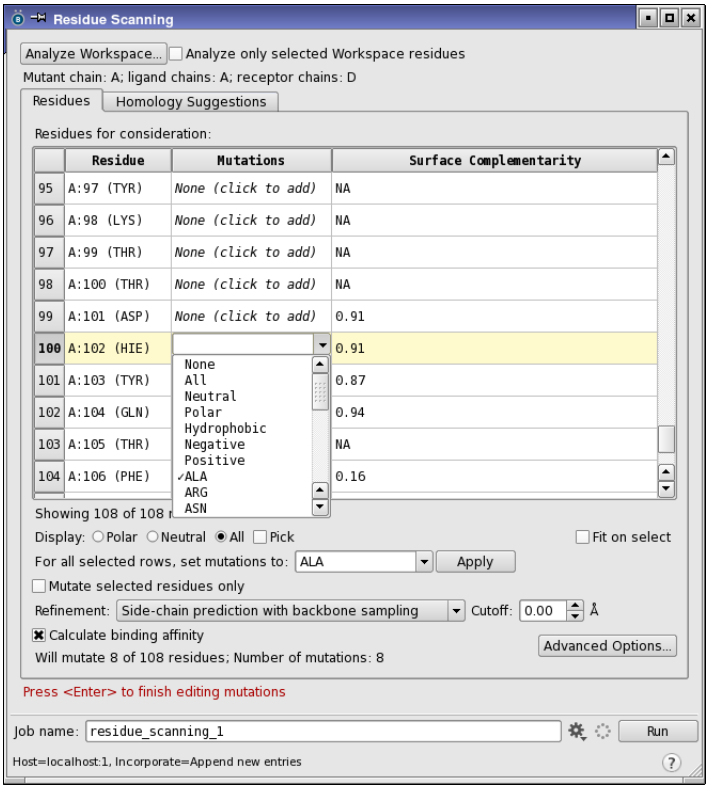

Figure 15.2. Option menu for defining mutations for a single residue.

The residue types include groups such as Neutral, Polar, Hydrophobic, Positive, and Negative. You can choose more than one residue or group from the list, and the new residues are added to the list. Each residue type that you select is checked on the list.

3. Click Apply.

The mutation list is set to the residues that you chose from the option menu. The mutations are listed in the Mutations column of the table.

A text message is displayed at the bottom of the tab that reads Will mutate *M* of *N* residues; Number of mutations: *K*. The number of mutations is the total number of structures generated. Only one site is mutated at a time, so the number of structures is linear in both the number of sites mutated and the number of mutations at each site.

When you select rows in the table, the Workspace view zooms in to the residues you have selected, if Fit on select is checked. The selected residues are highlighted with green carbons, and the remaining residues are dimmed. (You can adjust how far the view zooms in by making a setting in the Preferences panel. Choose Edit → Settings → Preferences, select Fitting under Workspace in the tree on the left, then enter a value in the Fit margin text box.)

15.2.3 Setting Optimization Options for Affinity Maturation

In affinity maturation, a search is performed to find the optimum binding affinity between two sets of protein chains, or the maximum stability of the entire protein. You can choose the property by selecting Affinity or Stability from the Property to optimize options. If you chose not to divide the protein chains into two groups, only the stability option is available. Both properties are calculated with Prime MM-GBSA in implicit solvent.

There are two choices for the search:

- Monte Carlo optimization—Perform a Monte Carlo optimization of the chosen property. At each Monte Carlo step, a residue is mutated at random, chosen from the set of all mutations for all residues to mutate. Residues that have already been mutated can be mutated again, and mutated back to the original residue. The acceptance criterion for the step (or move) is based on the change in affinity or stability. The process is repeated until the maximum number of steps given in the Maximum steps text box is reached.

  You can control some aspects of the Monte Carlo procedure with settings in the Affinity Maturation - Advanced Options dialog box. The effective temperature for the Boltzmann probability used in the acceptance test can be set, and the seed for the random selection of the mutation. The search can be terminated if no accepted move is found in a specified number of moves, and you can reject steps whose stability or affinity changes by more than a specified amount when optimizing on the other property (the affinity or the stability, respectively).

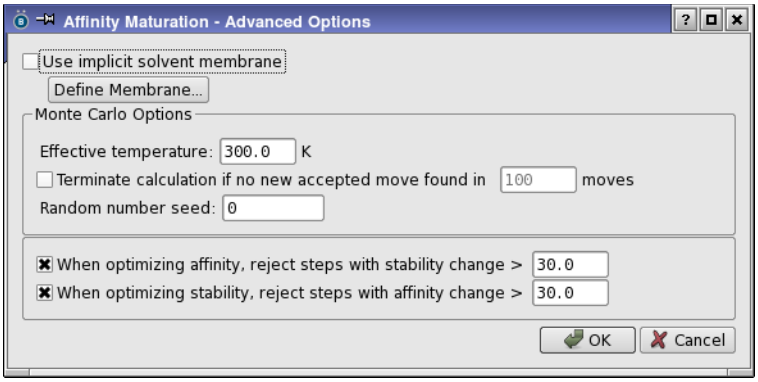

Figure 15.3. Affinity Maturation - Advanced Options dialog box

- Brute force exhaustive search—Perform all possible mutations of the residues to locate the structures with the best affinity or stability. You should keep the number of residues and mutations per residue small if you choose this option, to avoid combinatorial explosion. For example, mutating 5 residues with only one mutation per residue results in 31 possible mutated structures; with two mutations per residue the number of possible mutated structures is 242; with three mutations it is 1023; with four it is 3224; and with ten it is 161050.

You can set a limit on the number of mutations in any structure in the Maximum mutated residues per output structure text box. This restricts the search space, which is useful if you are making many mutations at a number of sites.

You can also specify the maximum number of output structures to return in the Maximum number of output structures text box. The structures that are returned are those with the best (most negative) value of the binding affinity or the stability.

15.2.4 Choosing Refinement Options

When the residues are mutated, the protein structure around the mutation site should be allowed to relax in response to the mutation. Relaxation is carried out based on the refinement method selected, which is described below.

The region around the mutation site that is relaxed during the calculations is defined by a cutoff distance, which you specify in the Cutoff text box. Any residue that has an atom within the specified distance of a hypothetical Arg residue at the mutation site is included in the refinement. A hypothetical Arg residue is used to ensure that the set of residues refined is identical, regardless of the initial or mutated residue identities. This ensures that comparisons of the properties of the mutated structures are not affected by the choice of residues that are relaxed.

Before the minimization of the region around the mutation site, a side-chain prediction of the mutated residues is performed, which does a thorough exploration of possible side-chain conformations. The method used for the side-chain prediction can be chosen from the Refinement option menu. The choices are:

- Side-chain prediction—perform only the side-chain prediction.
- Side-chain prediction with backbone minimization—perform the side-chain prediction with minimization of the backbone atoms.
- Side-chain prediction with Cbeta sampling—perform the side-chain prediction including sampling of the CA-CB orientation.
- Side-chain prediction with backbone sampling—perform the side-chain prediction including sampling of the backbone conformations.

In the Residue Scanning panel, you can calculate the binding affinity change on mutation for systems that consist of more than one chain, by selecting Calculate binding affinity. In the Affinity Maturation panel, the binding affinity is calculated if you divided the protein chains into two groups. Binding affinities are calculated with the Prime MM-GBSA technology.

If you want to use an implicit membrane model for the protein, you can set it up by clicking Advanced Options. Select Use implicit solvent membrane, and click Define Membrane to place the membrane on the protein. The membrane is modeled with a low-dielectric slab.

15.2.5 Running the Job

To run the job, you first set job parameters, and then submit the job to a host for execution. Click the Settings button to open the Job Settings dialog box. In this dialog box, you can choose whether to append the results to your project, or leave the results on disk in the current directory, where you can examine them later. You can also choose a job name, which is used to name the files associated with the job. When you have made your choices, click Save and Run in the dialog box to start the job.

If you run the job on a multiprocessor host, you can divide the job into subjobs and distribute them over multiple processors. The minimum work a subjob can do is to mutate one residue to another residue, so you should not specify more subjobs than there are mutations. For optimal load balancing, the number of subjobs should be a few times the number of processors.

Monte Carlo affinity maturation jobs perform a random walk with a different random seed in each subjob if the job is distributed over multiple processors. The results of each walk are collated at the end to select the structures with the best values of the stability or affinity.

A status bar showing the progress of the job is displayed above the Start button. The job takes several minutes per residue mutation to run, depending on the refinement options.

15.3 Using Homologs for Identifying Mutations

In addition to manual selection of residues, you can use homology to suggest residues to mutate and residues to mutate to, on the basis of variability or conservation of residues across the set of homologs, or on structural proximity or properties. You can do this in the Homology Suggestions tab.

Homology suggestions are generated for a single chain. For residue scanning, this is the chain you selected for mutation. For affinity maturation, you can choose the chain from the Chain option menu, and you can generate suggestions for each chain in turn.

15.3.1 Obtaining Homologs from a BLAST Search

If you don't have homologs for your protein, you can run a BLAST search to find homologs.

1. Click Run a Blast Search.

   The BLAST Search Settings dialog box opens, so you can change settings if you want.

2. Click Start Job in this dialog box, after changing any settings.

   The job starts, and its log file is displayed in the Job Progress dialog box.

At the end of the job, the BLAST Search Results dialog box opens, so that you can choose how many of the homologs you want to use. You can do one of the following to select homologs:

- Select the homologs in the table (with shift-click and control-click).
- Enter the number of the top homologs in the text box at the bottom of the panel and click Select Top.

When you have selected the homologs, click Incorporate Selected Rows to add the structures to the sequence viewer.

15.3.2 Importing Homologs

If you already have a set of homologs for a structure that you want to use to identify potential mutation sites, or if you have other structures that you want to use as homologs, you can import them.

- If you have a set of homologs in a file, click the Import button and choose Browse for file. A file selector opens so you can locate and import the file.
- If you want to manually import structures from the PDB, click Import and choose From PDB ID, then specify the PDB ID in the dialog box that opens.
- If the structures are in the Workspace, click Import and choose From Workspace.

15.3.3 Aligning Homologs

To make use of the homologs, they must be aligned to the parent (query), if they are not already aligned. You can run a job to align the homologs, or you can do a manual alignment, or both.

To run the alignment job, click Align Homologs, or click the Multiple Alignment toolbar button.

The alignment, which uses ClustalW, is usually very quick. A dialog box may be displayed briefly before the results are incorporated. The alignment adds gaps in the sequence viewer as necessary.

To align sequences manually, you can drag residues to the right or the left, to fill or create gaps. If you have already created gaps that you want to preserve during another alignment, you can

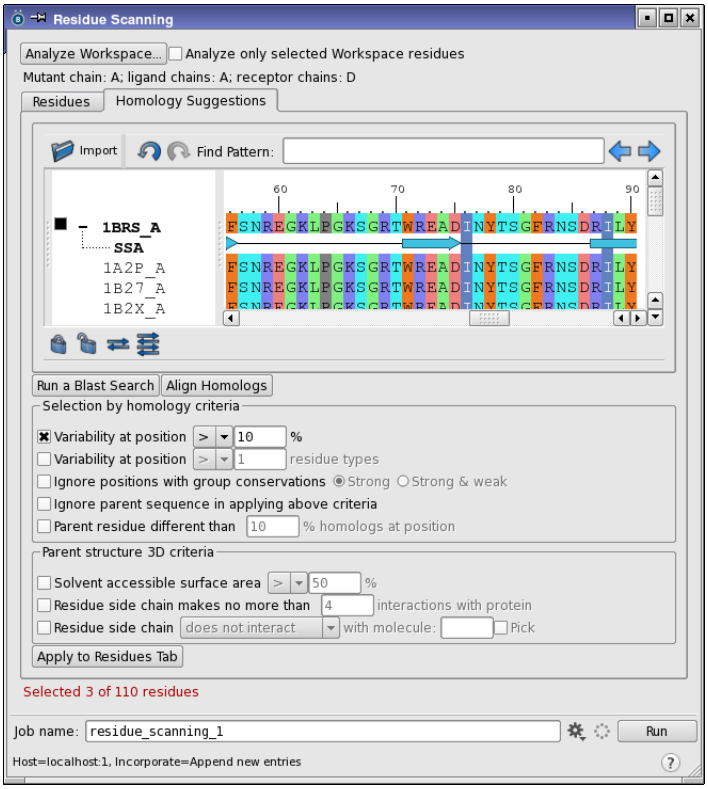

Figure 15.4. The Residue Scanning panel, Homology Suggestions tab.

click the Lock Gaps toolbar button. The gap symbol changes to - to indicate that the gap is locked. To unlock the gaps again, click the Unlock Gaps toolbar button.

15.3.4 Selecting Residues by Homology

When the sequences are aligned, residues can be selected on the basis of the of the alignment. To choose residues to mutate, you may want to select residues at positions where the variability among the aligned homologs is high, or residues that vary in residue type, or residues that differ from most of the homologs. The Selection by homology criteria section has a number of options that you can set to apply selection filters on the basis of the amount of variation. Each of the options that you select is applied, so the residues must meet all specified criteria.

- Variability at position >/< *N* %—Select residues based on the percentage variability at the residue position. You can apply a minimum or a maximum variability by choosing from the option menu, and specify the percentage threshold in the text box.
- Variability at position >/< *N* residue types—Select residues based on the residue type variability at the residue position. Choose whether to apply a minimum or a maximum variability from the option menu, and specify the threshold for the number of residue types that can vary in the text box.
- Ignore positions with group conservations—Ignore (do not select) residues that are strongly conserved or that are either strongly or weakly conserved. Select the appropriate option for strong conservation or both strong and weak conservation.

  Strong conservation means that all residues at a particular position are in one of the following groups: STA, NEQK, NHQK, NDEQ, QHRK, MILV, MILF, HY, FYW. Weak conservation means that all residues at a particular position are in one of the following groups: CSA, ATV, SAG, STNK, STPA, SGND, SNDEQK, NDEQHK, NEQHRK, FVLIM, HFY. These definitions are those used by ClustalW.
- Ignore parent sequence in applying above criteria—Ignore the parent (query) sequence when applying the variability or conservation criteria. The selection is then based on the variability in the homologs only.
- Parent residue different than *N* % homologs at position—Select residues for which the parent residue is different from more than the specified percentage of homologs at the residue position.

When residues are selected on the basis of homology, a default set of mutations is also defined. The mutations for a given residue include all the variants found at that residue position.

15.3.5 Selecting Residues by Structural Attributes

Residue selection filters can also be applied on the basis of solvent-accessible surface area and contact with other parts of the protein or with a ligand or cofactor. To apply these criteria to residue selection, select one or more of the following options:

- Solvent accessible surface area—Select this option to select residues by their solvent-accessible surface area (SASA) relative to an isolated residue of the same type, and set a threshold for the maximum or minimum allowed relative SASA. This option is useful for locating surface or buried residues.
- Residue side chain makes no more than *N* interactions with protein—Select this option to filter out residues whose side chains have multiple interactions with other protein residues, and set the maximum number of residues with which the side chain has interactions.
- Residue side chain interacts/does not interact with molecule *N*—Select this option to select residues by their interaction with a particular molecule. Choose whether to allow or disallow the interaction from the option menu, and specify the molecule in the text box, or select Pick and pick the molecule in the Workspace. This is useful for mutating residues near a ligand, or for mutating residues at protein-protein interfaces, for example.

Interactions are determined by a distance cutoff: any residue that has atoms within 4 Å of the side chain is considered to interact.

15.3.6 Making the Selection

When you have set up all the criteria for residue selection, click Apply to Residues Tab. The rows are selected in the table in the Residues tab, for the residues that meet all the criteria. For each selected residue, a set of mutations is defined that correspond to the residues observed for that position in the homologs that were used in the Homology Suggestion tab. You can then use the other tools in the Residues tab to modify any of the selections before running the job.

15.4 Examining the Mutation Results

When the job finishes, the Residue Scanning Viewer or Affinity Maturation Viewer panel opens, displaying changes in properties for each mutation, and a graph of one of the properties against the row number.

If you want to examine results from a job that was completed earlier, you can open this panel by choosing Tasks → Residue Scanning → View Results or Tasks → Affinity Maturation → View Results. You can then click Import to locate the Maestro file (`.maegz`) that contains your results and import it into the current project.

The results are listed in the Mutations table, with one mutated structure per row. The mutations in each structure are identified by the residue positions and the original and mutated residue names. For residue scanning, there is only one mutation per structure, but for affinity maturation, there may be multiple mutations. The properties that were calculated for each mutant are listed in the table. These properties are described briefly in Table 15.1. Some of these properties are described in more detail in the sections below. In addition to the properties listed, the change in the Prime energy properties are also available. All properties are calculated after the refinement is performed, and so include relaxation of the protein after mutation.

For affinity maturation, residue-level properties are summed, so they represent total changes due to the mutations. Dividing by the number of residues would give the average change in these properties. The properties for the individual residues are not reported. You can also create a LOGO plot from the affinity maturation results by clicking Create LOGO plot. The plot is written to a `.png` file and the data to a `.csv` file with the job name as the base name.

You can sort the table columns by clicking on the column headings. You can plot any of these properties against the mutation (table row) by choosing the property from the Graph property option menu. If you want to export the table data as a CSV file, click Export, and navigate to a location and name the file.

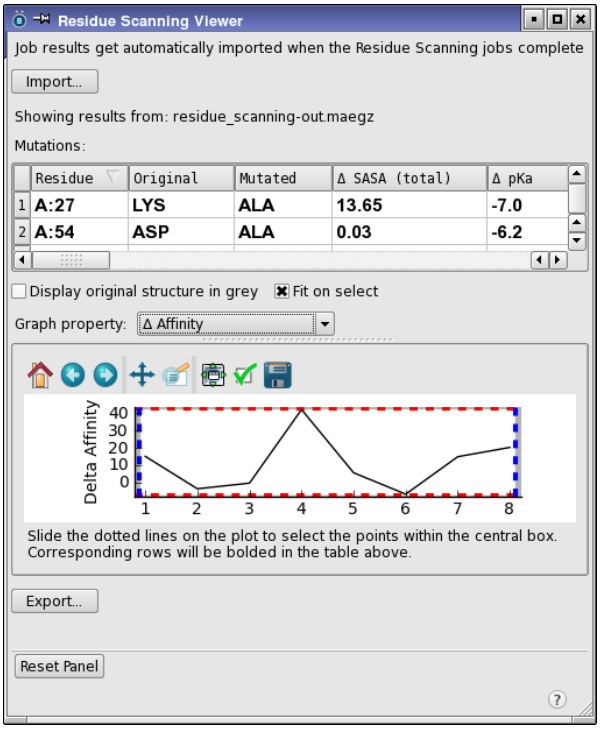

Figure 15.5. The Residue Scanning Viewer panel.

Table 15.1. Mutation properties.

Property	Description
Δ SASA (total)	Change in total surface area due to the mutation.
Δ SASA (nonpolar)	Change in surface area of nonpolar atoms due to the mutation
Δ SASA (polar)	Change in surface area of polar atoms due to the mutation
Δ pKa	Change in pK_a of the mutated residue, calculated using PROPKA [1].
Δ Affinity	Change in binding affinity of the mutated protein (and any other specified chains), treated as the ligand, to the rest of the system, treated as the receptor. A negative value means that the mutant binds better than the native protein. The calculations are carried out with Prime with an implicit solvent term (see Section 15.4.1 on page 110 for details).
Δ Hydropathy	Change in hydrophobic or hydrophilic nature of the mutated residue, as defined on the Kyte-Doolittle scale [6] - see, for example, http://en.wikipedia.org/wiki/Hydrophobicity_scales. A positive values indicates a more hydrophobic residue and a negative value indicates a less hydrophobic residue in the mutant.
Δ Total rotatable bonds	Change in the total number of rotatable bonds.
Δ Potential Stability	Change in the stability of the protein due to the mutation, calculated with just the potential energy terms (no solvation). The stability is defined as the difference in energy between the folded and unfolded states (see Section 15.4.2 on page 110). A negative value of the stability means that the mutant is more stable than the native protein.
Δ Stability (solvated)	Change in the stability of the protein due to the mutation, calculated using the Prime energy function with an implicit solvent term. The stability is defined as the difference in free energy between the folded state and the unfolded state (see Section 15.4.2 on page 110). A negative value of the stability means that the mutant is more stable than the native protein.
Δ vdW Surface Complementarity	Change in the van der Waals surface complementarity of residues at chain interfaces due to the mutation.

You can select a region in the graph using the horizontal and vertical dashed lines, which can be dragged to create the selection. The rows corresponding to the selected region of the graph are highlighted in the table above, and the residues are highlighted in the Workspace.

If you select a table row, the view zooms in to the mutated residue. To display the original structure, select Display original structure in grey. The parent protein is displayed and colored grey. You can then see how the mutation is positioned in relation to the original residue.

15.4.1 Binding Affinity Prediction

The change in the binding affinity of the protein due to the mutation is calculated from a thermodynamic cycle, which can be represented as follows:

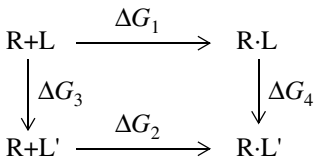

where R is the receptor, L is the ligand in the parent, and L' is the mutated ligand. R+L and R+L' represent the separated receptor and ligand. R·L and R·L' represent the receptor bound to the ligand. The change in binding affinity is

$$\Delta\Delta G(\text{bind}) = \Delta G_2 - \Delta G_1 = \Delta G_4 - \Delta G_3$$

Experiment measures ΔG_1 and ΔG_2, but it is ΔG_3 and ΔG_4 that are calculated, to optimize the cancellation of error in the computational models. The calculations are done with Prime MM-GBSA, which uses an implicit (continuum) solvation model. A negative value indicates that the mutant binds better than the parent protein.

15.4.2 Stability Prediction

The stability of the protein due to the mutation is calculated from a thermodynamic cycle, which can be represented as follows:

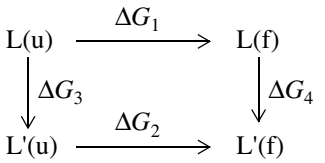

where L(u) is the unfolded parent ligand, L(f) is the folded parent ligand, L'(u) is the unfolded mutated ligand, and L'(f) is the folded mutated ligand. The change in stability is

$$\Delta\Delta G(\text{stability}) = \Delta G_2 - \Delta G_1 = \Delta G_4 - \Delta G_3$$

Experiment measures ΔG_1 and ΔG_2, but it is ΔG_3 and ΔG_4 that are calculated, to optimize the cancellation of error in the computational models. For the purpose of the model, the unfolded ligand is represented as a tripeptide, Gly-X-Gly, where X is the residue that is mutated. The assumption is that the remaining interactions in the unfolded state are negligible. The calculations are done with Prime MM-GBSA, which uses an implicit (continuum) solvation model.

Locating Possible Mutations for Disulfide Bridges

Disulfide bridges between cysteine residues add to the stability of a protein structure. Mutating residues to form or break disulfide bridges offers a way of controlling the stability of a protein. This chapter describes how to run a cysteine mutation calculation to locate and rank possible disulfide bridges. The calculation is set up and run in the Cysteine Mutation panel, in the Run tab, and the results are presented in the Results tab.

To open the Cysteine Mutation panel, choose Tasks → Cysteine Mutation.

16.1 Selecting and Analyzing the Protein

The Cysteine Mutation panel can analyze a single protein for potential mutations, such as a structure from the PDB, or it can analyze a molecular dynamics trajectory, to find frames in which suitable residues come close enough to form disulfide bridges.

If you want to analyze a single protein, first display it in the Workspace. The protein must be one that has been prepared for use in modeling. If it has not been prepared, we recommend that you prepare it with the Protein Preparation Wizard (on the Tools menu and the Tasks menu). Details of preparing a protein can be found in the *Protein Preparation Guide*.

The protein must be analyzed to locate possible residue pairs that could be mutated to cysteines, or to locate disulfide bridges that could be broken by mutation to other residues, which you do by clicking Analyze Workspace. Instead of analyzing the entire protein, you can analyze the Workspace selection. To do this, select the desired residues in the Workspace structure, select Analyze only selected Workspace residues in the Cysteine Mutation panel, then click Analyze Workspace.

If your structure has not been prepared for modeling, you are asked if you want to use the Protein Preparation Wizard to prepare the structure. You cannot proceed if you do not have a protein that is an all-atom, 3D structure with bonding information. After preparing the protein, display it in the Workspace again and click Analyze Workspace.

If you want to analyze a molecular dynamics trajectory, click Analyze MD Trajectory. A file selector opens, in which you can locate a Desmond MD simulation results file (`-out.cms`). After selecting the file, you are prompted to specify the interval at which the analysis is performed on the trajectory, as a number of steps. The analysis takes some time, and is run under job control. Running a Desmond MD simulation requires prior protein preparation, so no further preparation is necessary.

When the structure has been analyzed, the Residue pairs for mutation table is filled in with all the residue pairs that meet the criteria for forming or breaking a disulfide bridge. The criterion for identifying potential cysteine pairs is a Cβ–Cβ distance between the residues that is less than the distance specified in the panel. For Gly, the distance is taken from the alpha hydrogen.

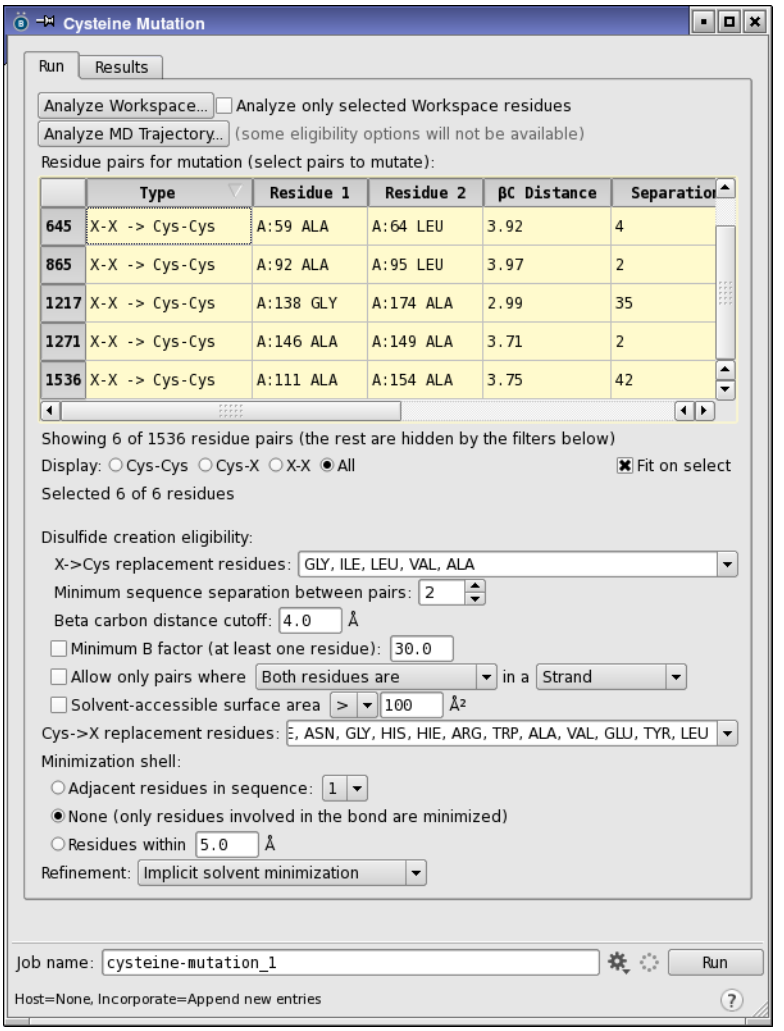

Figure 16.1. The Run tab of the Cysteine Mutation panel.

16.2 Choosing Residue Pairs to Mutate

The Residue pairs for mutation table lists the pairs of residues that have the potential to form or break disulfide bonds. The table rows that are shown are controlled by the options below the table. The table columns include a range of relevant properties, which are described in Table 16.1. You can sort the table by the values in a column by clicking on the column heading.

Table 16.1. Columns of the Residue pairs for mutation table.

Column	Description
(Index)	The first column contains the index of the residue pair. When the table is filtered to show only certain residue pairs, this index remains the same (i.e. it is not a table row number).
Type	Mutation type, which can be one of the following: X-X -> S-S: Mutation of two residues to Cys with formation of a disulfide bond. S-X -> S-S: Mutation of one residue to Cys with formation of a disulfide bond to a nearby cysteine S-S -> S-X: Mutation of one Cys residue of a bonded pair to break a disulfide bond.
Residue 1	Identity of the first residue in the pair, given by the chain letter, the residue number and insertion code, and the 3-letter residue name. For Cys-Cys pairs, this residue is the residue that is mutated, so the pair is listed twice, in opposite order, to allow selection of only one of the pair to mutate.
Residue 2	Identity of the second residue in the pair, given by the chain letter, the residue number and insertion code, and the 3-letter residue name.
β-carbon Distance	Distance in angstroms between the beta carbons (CB) of the two residues. In the case of Gly, the distance is taken from the alpha hydrogen (HA) that would be replaced by the beta carbon of the Cys.
Separation	Sequence separation between the residues in the pair, defined as the number of residues between the two residues. Displayed as N/A if the residues are in different chains.
SASA	Combined solvent-accessible surface area of the two residues in the pair.
Sec. Structure	Secondary structure elements of the two residues in the pair (helix, strand, etc.). If both have the same secondary structure element, it is only given once, otherwise the two elements are given in the form element1/element2.
B Factor	Temperature factors (crystallographic B factors) of the two residues in the pair, represented as B(residue 1)/B(residue 2).
Frame	Trajectory frame, if an analysis was performed on a trajectory. Each pair can come from a different frame, and the structure in that frame is the one that is mutated.

There are several ways to control what is shown in the table.

a. Use the Display options:

   - Cys-Cys—Show only cysteine-cysteine pairs, for S-S -> S-X mutations. The first residue listed is the one mutated, so a given pair is listed twice, in opposite order, to allow selection of either residue of the pair for mutation.
   - Cys-X—Show only pairs with one cysteine, for S-X -> S-S mutations.
   - X-X—Show only pairs of non-cysteine residues, for X-X -> S-S mutations.
   - All—Show all pairs.

b. Use the X->Cys replacement residues option menu to filter on the residues that will be replaced with Cys.

   The option menu contains individual residues, which you can select independently, an All option to select all residues (except Pro), a None option to clear the list, and a Conservative option to select conserved residues (GLY, ILE, LEU, VAL, ALA). The residues that you choose are displayed in the main part of the option menu; the complete list is shown in a tooltip if it is too long. The table is updated for each selection that you make.

c. Use the Beta carbon distance cutoff text box

   You can specify the maximum allowed distance between the beta carbons of the residues in a pair. This distance is used to filter the table to show only residues with a smaller distance. For Gly, the distance is taken from the alpha hydrogen.

d. Use the Minimum B factor text box

   Specify the minimum B factor that at least one residue in the pair must have.

e. Restrict the secondary structure elements that the residues belong to, by selecting Allow only pairs where, and choosing from the menus

   The first option menu allows you to include or exclude secondary structure elements for one or both residues, and the second option menu allows you to choose the secondary structure elements.

f. Use the Solvent accessible surface area option, menu, and text box

   Specify the minimum or maximum combined SASA for the residue pair.

These display options allow you to restrict the list of residue pairs to those that are of interest, so that it is easier to select the pairs in the table that you want to mutate. The job mutates only the selected pairs in the table, so you must make a selection before you run the job.

Each selected pair is mutated independently: there are no simultaneous mutations.

If you have Cys-Cys pairs whose bond you want to break by mutating one of the cysteines to another residue, you must also select the mutations, from the Cys -> X replacement residues option menu. The option menu works in the same way as the X -> Cys replacement residues option menu, described above.

16.3 Relaxing the Structure Around the Mutation Site

When a pair of residues is mutated, the side-chain conformations are sampled to find the best geometry. The residues are then minimized, to relieve strain. You can include additional residues around the mutation site in the minimization, to extend the relaxation of the protein. There are three options for selecting the residues, listed under Minimization shell:

- Residues within *N* Å—Optimize all residues that have atoms within the specified distance of the mutated residue pair.
- Adjacent residues in sequence—Optimize the *N* residues next in the sequence on either side of the residues in the pair, where *N* is selected from the option menu.
- None—Do not optimize any residues but the two residues in the pair.

You can also choose whether to run the minimization in the gas phase (the fastest option), or use an implicit solvent model, or not to minimize at all. If the solvent-accessible surface area of the minimization region is negligible, the implicit solvent model may not be of any value; if the SASA is large enough, the implicit solvent model should probably be used.

16.4 Running the Job

When you have selected the residue pairs that you want to mutate, and set any options for relaxing the structure around the residue pairs, enter a name in the Job name text box, and click Run.

If you want to choose whether to incorporate the results into the project, click the Settings button. The Job Settings dialog box opens. There are two output options:

- Append new entries—Append the mutated structures to the project.
- Do not incorporate—Leave the mutated structures in the working directory.

The Maestro file containing the structures is copied to the working directory when the job finishes, regardless of the option. If you choose not to incorporate the results into the project, you can always do so later by importing the output file. When you have chosen the output option and named the job, click Save and Run.

The job is run locally, and its progress is displayed in a status bar at the bottom of the panel.

16.5 Examining the Results

The results of a cysteine mutation job are automatically displayed in the Results tab if you chose to append new entries in the Start dialog box. To view results of any other cysteine mutation job, click Load Results from Previous Run. This button opens a file selector, in which you can browse to the output Maestro file (`-out.maegz`) and load it.

The results of the mutation job are displayed in the Mutations table. The table columns are described in Table 16.2.

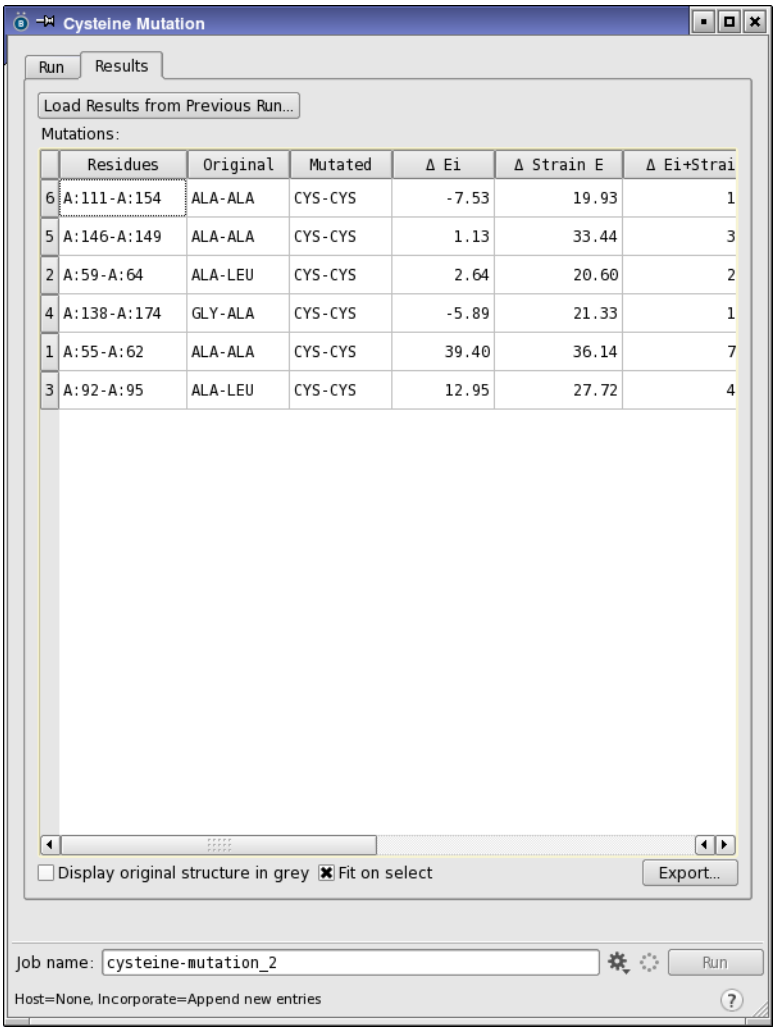

Figure 16.2. The Results tab of the Cysteine Mutation panel.

Selecting a row in the table zooms in on the mutated pair in the Workspace, if you have Fit on select selected. The mutated residues are displayed in ball-and-stick representation, and the rest of the structure uses a darker color scheme. If you want to see the original residues as well, select Display original structure in grey.

Table 16.2. Columns of the Mutations table.

Column	Description
Residues	Chain name, residue number and insertion code of both residues in the mutated pair.
Original	3-letter names of original residues in the pair.
Mutated	3-letter names of mutated residues.
Δ Ei	Change in interaction energy between the residue pair and the rest of the protein on mutation.
Δ Strain E	Change in strain energy on mutation. The strain energy is the difference in internal energy between the state of the residue pair in the protein and the relaxed state of each residue or of the disulfide in the gas phase.
Δ Ei+Strain E	Change in the sum of interaction energy and strain energy on mutation, equal to Δ Ei + Δ Strain E.
Pre-min Score	Geometric energy score of the mutated protein prior to minimization. The score is calculated using an empirical function that is derived from the distributions of the internal coordinates of cysteine disulfides in the PDB. Geometries that are seen in the PDB for disulfides yield lower scores.
Weighted Score	Weighted sum of change in interaction energy, change in strain energy, pre-minimization score and post-minimization score. The last of these uses the same method as the pre-minimization score and is reported in the Project Table, but is not reported here. The score includes a shift for mutations for which any of the four components of the weighted score is larger than a threshold for that component. Sorting the weighted score places the mutations that pass all threshold tests first, followed by all the others.

16.6 Workflow Summary

1. Display the protein you want to analyze in the Workspace. You should ensure that it is properly prepared (with the Protein Preparation Wizard).
2. If you want to analyze only part of the protein, select the residues that you want to analyze, and select Analyze only selected Workspace residues.
3. Click Analyze Workspace or Analyze MD Trajectory.

   When the analysis finishes, the Residue pairs for mutation table is populated with the residue pairs that were found.
4. Filter the list to show only the mutations you are interested in:
5. Use the display options to display the mutation types that you are interested in.
6. Use the menus below the table to select the residues that you want to mutate to or from. The table is updated to show only those residues.
7. Use the cutoff to filter out pairs for which the distance is too large to form a disulfide bond.
8. Select the pairs that you want to mutate in the table. Each pair is mutated independently to produce a structure and an energy evaluation.
9. Decide whether you want to optimize residues near to the mutated residues by selecting from the Gas phase optimization shell options. The mutated residues are always optimized.
10. Enter a name in the Job name text box and click Run to run the mutation job. To decide how to handle the output, click the Settings button instead.
11. When the job finishes, go to the Results tab. If the results are not listed in the Mutations table, click Load Results from Previous Run, navigate to and open the output Maestro file from the job.

Antibody Modeling

The main task of antibody modeling in BioLuminate is to construct a homology model based on a database of antibody templates. Once the model is constructed, it can be humanized if required. A database of antibody templates is provided with BioLuminate, based on a new analysis of antibodies in the PDB from 2010 [2], which you can modify or add to, or create your own database. Antibody-antigen docking is covered in Section 11.5 on page 75.

17.1 Modeling an Antibody Structure

Modeling of the Fv region of an antibody involves prediction of both the framework (FR) region and the variable loop (CDR) region. These regions are identified automatically and their structure predicted by homology, based on known antibody structures from the PDB, or from your own database of antibodies. You can also use input coordinates for some parts of the structure and predict the rest. Finally, the H3 loop can be refined after the initial structure is generated. You can also model the entire antibody including the constant region (Fc).

The modeling is performed using the Antibody Prediction panel, which you open by choosing Tasks → Antibody Modeling → Prediction in the main window.

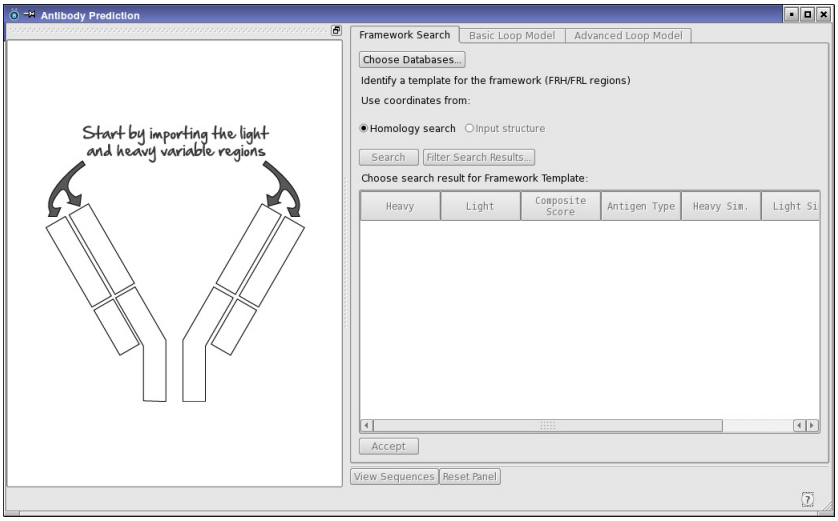

Figure 17.1. The Antibody Prediction panel, initial view.

17.1.1 Importing the Antibody Sequence

The first step in the modeling process is to import the antibody sequence, and if you want to re-predict only a part of the antibody, the existing structure.

The left part of the panel displays a diagram of the Fv region of an antibody. Clicking on the light variable or heavy variable region in the diagram displays a menu, from which you can choose the source of the sequence for this region. The choices are:

- Browse for File—Open a file browser in which you can navigate to the desired location and select the file that contains the sequence.
- From Workspace—Use the sequence for the structure that is displayed in the Workspace. Only one structure must be displayed.
- From Selected Entries in the Project Table—Use the sequence from the entry that is selected in the Project Table. Only one entry must be selected.
- From PDB ID—Use the sequence from the specified PDB ID. Opens the Enter PDB ID dialog box, in which you can enter the PDB ID of the sequence. The sequence is retrieved from a local copy of the PDB if it is available, or from the RCSB web site, depending on the preference set for PDB retrieval.
- Enter/Paste New Sequence—Type or paste in the sequence. Opens the Sequence Editor dialog box, in which you can name the sequence and type it or paste it in, as a string of single-letter codes.

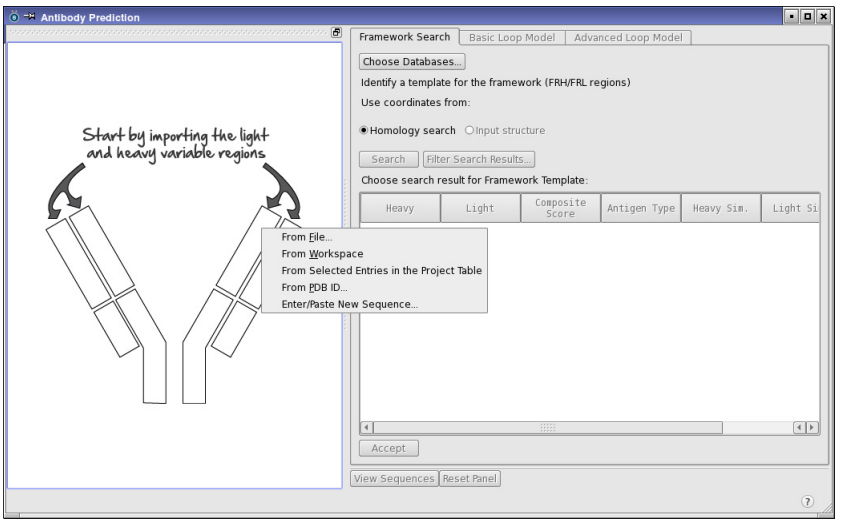

Figure 17.2. The Antibody Prediction panel showing the import menu.

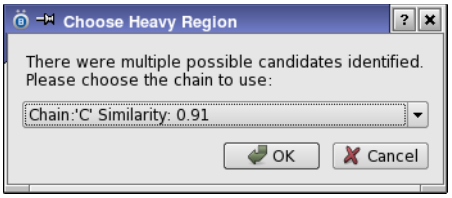

Figure 17.3. The Choose Heavy Region dialog box

If you intend to model only part of an antibody and use the input structure for the rest, you should ensure that you import a sequence that has an associated structure. This means that you can choose any of the menu items except the last.

When the protein is read in, it is analyzed to find the chains and the loops. If there is more than one possible chain that could be used (for example, in a dimer), a dialog box opens, in which you can choose one of the chains. When the analysis finishes and you have chosen a chain if requested, the region in the diagram is colored to indicate that the chain is assigned. When both chains have been chosen, the text prompting you to import the two chains is no longer displayed.

You can view the sequences in a sequence viewer at any time by clicking View Sequences.

17.1.2 Choosing a Database

The model of the antibody is built using a database of antibodies, each of which has been analyzed and curated in terms of framework and hypervariable loop regions. This database is used to model antibody structures from a sequence. A curated database obtained from the PDB is provided with the software, but you can add your own structures to customize it, or build your own databases—see Section 17.4 on page 142.

You can choose one or more databases to use when modeling an antibody. The default is the database in the installation. If you want to select the databases, click Choose Databases, to open the Choose Database dialog box. This dialog box has a table of databases that you can choose from. You can do the following:

- Add a database to the list, by clicking Add Database and navigating to the database in the file selector that opens.
- Select a database to use, by checking the check box in the Active column.
- Remove a database from the list, by clicking the button in the Actions column.

When you have finished modifying the list and choosing the active databases, click OK.

17.1.3 Selecting the Coordinates for the Framework Region

You can choose between two sources for the coordinates of the framework region: the input structure, or a homology model.

If you want to use the coordinates from the input structure, select Input structure. The sequences that you imported for the light and heavy regions must of course be associated with a structure, and an error message is posted if there is no structure available.

If you want to use a homology model, select Homology search. To run the search, click Search. When the search finishes, the table in the lower part of the Framework Search tab is filled in with the results, in order of their score. The table columns are described in Table 17.1. You can select a single result in the table to use as the template.

Table 17.1. Description of search results table columns.

Column	Description
Heavy	PDB ID of the homolog for the heavy framework region
Light	PDB ID of the homolog for the light framework region
Composite Score	Average of the Heavy Sim. and Light Sim. scores. This score is used to order the homologs in the table.
Antigen Type	If the template contains an antigen, its type and chain length are shown in this column. If it does not contain an antigen, the column shows None.
Heavy Sim.	Sequence similarity of the entire variable domain sequence (framework and CDR) for the heavy chain. The similarity is the number of matching residues divided by the total number of residues, where "matching" means that the two residues have a positive score in the BLOSUM62 matrix.
Light Sim.	Sequence similarity of the entire variable domain sequence (framework and CDR) for the light chain. The similarity is the number of matching residues divided by the total number of residues, where "matching" means that the two residues have a positive score in the BLOSUM62 matrix.

When you have selected a template or chosen to use input coordinates, click Accept to accept the choice and move on to the next stage. The Basic Loop Model tab is displayed automatically, after a short pause.

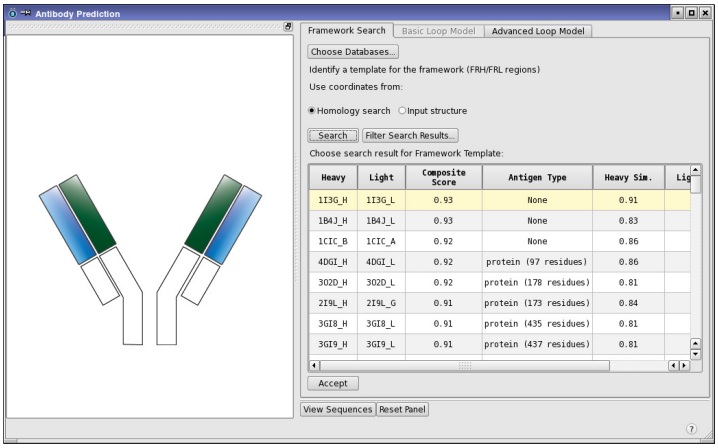

Figure 17.4. The Antibody Prediction panel after searching for templates.

17.1.4 Filtering the Database by Property

By default, the homology search scans all of the structures in the databases you have chosen. If you want to restrict the search to only the structures that have certain properties, you can set up a filter on these properties before conducting the search. To set up the filter, click Filter Search Results, which opens the Filter Search Results dialog box.

First, choose a property from the Available properties list at the top of the panel. The list shows the property name, the family (category) it belongs to, and the range of values, for numeric properties. You can limit the list to a particular property family by choosing from the Show family option menu. If you type in the Property text box, a completion list is displayed below it, from which you can choose a property.

Once you have chosen a property, it is displayed in the Property text box. You can then use the text boxes and menus to the right to define the restrictions on the values of this property. Click Add to add the filter to the Filtering definitions and criteria list. You can add multiple criteria or definitions, and each of them is applied to the databases. The number of structures in the database that match the filters is reported at the end of the list. The search is case-insensitive, so for example 1FSK matches 1fsk. To do case-sensitive searches, select Case-sensitive.

If you want to filter on multiple properties, you can choose another property, set up the restrictions on its value, and click Add to add the filter to the list. The filters are cumulative (implicit AND), so the resulting structures are those that match both filters.

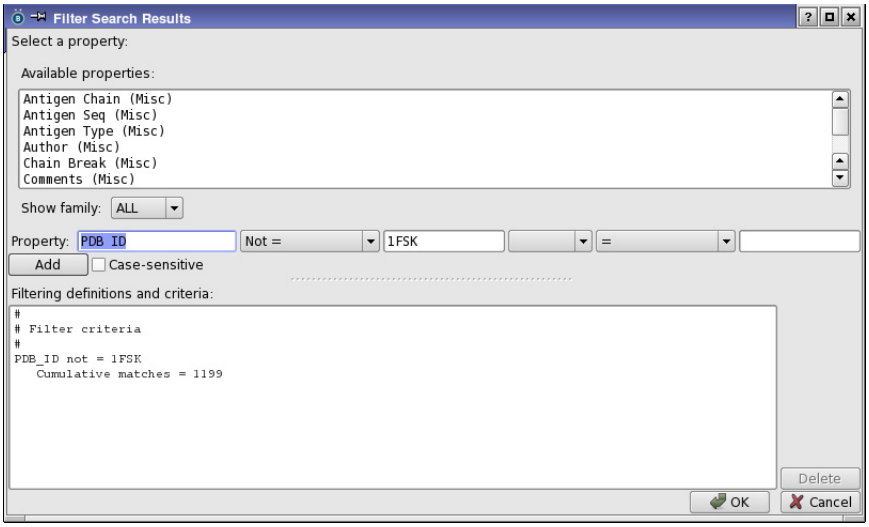

Figure 17.5. The Filter Search Results dialog box.

As an example, to include structures with properties in a specified range (all values from a minimum to a maximum inclusive):

1. Choose >= from the first menu.
2. Enter the minimum in the next text box.
3. Choose AND from the second menu.
4. Choose <= from the third menu
5. Enter the maximum value in the final text box.

As another example, to filter our structures for which a property has values in a given range:

1. Choose < from the first menu
2. Enter the minimum in the next text box
3. Choose OR from the second menu
4. Choose > from the third menu
5. Enter the maximum value in the final text box.

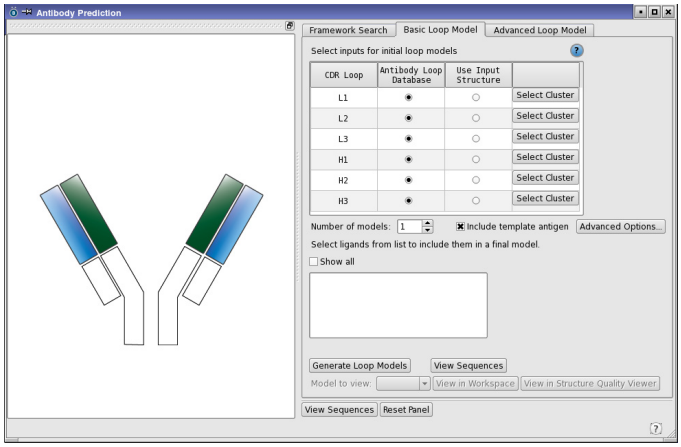

Figure 17.6. The Antibody Prediction panel, Basic Loop Model tab.

17.1.5 Generating the Loop Model

With the framework region defined, you can now generate a model for the six loops. As for the framework region, you have the choice of using input coordinates for a loop or predicting the loop from the database.

The six CDR loops are listed in a table that provides buttons for selecting the source of input coordinates. To choose the source, click the appropriate button. You can only choose to use the input structure if the framework also came from the input structure. By default all loops are predicted from the database. Accepting the default loop assignments is usually the best option.

If you want to choose individual clusters for one or more loops, click Select Cluster to open the Loop Clusters dialog box, which shows you the clusters for each loop, and allows you to select a cluster to use in the model. It also allows you to set a similarity cutoff for the loop to be used in the model. The procedure is described in more detail in Section 17.1.6 on page 126. Choosing a cluster automatically selects use of the antibody loop database for the loop.

If you want to include the template antigen in the model, select Include template antigen. This option is only available if the template has an antigen, which is indicated in the list of templates in the Framework Search tab.

Likewise, if you want to include ligands, cofactors, or water in the model, select them in the ligands list. Only ligands are displayed by default. To display the cofactors and water, select Show all above the list.

If you want to model the entire antibody including the constant region, click Advanced Options, select Include Fc region in the model, choose a template from the table, and click OK. There are very few templates, so you should be aware that the quality of the results in this region might be low. The model of the Fc region is built even if the sequence for this region is incomplete or missing, by using the template sequence.

Otherwise, the selection of the loop is done by sorting the clusters by the size of the cluster, then locating the largest cluster in the list that has a member whose loop similarity to that of the query is greater than the loop similarity cutoff. The cluster member with the greatest similarity to the query is the one that is used to build the loop.

You can build more than one model for the structure, by setting the desired number in the Number of models text box. If more than one model is requested, a series of diverse models is returned. The first model returned is usually the most likely to be correct.

To generate the loop model or models, click Generate Loop Models. To assess the quality of a generated model, you can examine it in the Workspace (click View in Workspace), or analyze it in the Structure Quality Viewer (click View in Structure Quality Viewer). If you have multiple models, you can choose them from the Model to view option menu.

When you view a model in the Workspace, it is colored by residue with the following color scheme:

- Blue—residues for which the full residue conformation was copied from the template.
- Cyan—residues for which the residue backbone conformation was copied from the template, and the side chain was modeled (because there was a residue mutation in the template relative to the query).
- Red—residues for which both the backbone and the side chain were modeled.
- Maroon—residues in the CDR loops.

When you view the models in the Protein Structure Quality Viewer, all models are listed in the table at the top of the viewer, and the structures are colored in the Workspace according to the Ramachandran plot regions.

17.1.6 Selecting Clusters for a Loop

If you want to manually select the loops that are used to build the model, you can do so in the Loop Clusters dialog box, which you open by clicking View Clusters in the Basic Loop Model tab. In this dialog box you can examine the loop clusters for each of the CDR regions, filtered by similarity, and select a cluster for each loop. You can also set the minimum similarity that is used in selecting a cluster, both manually and in the default automatic procedure.

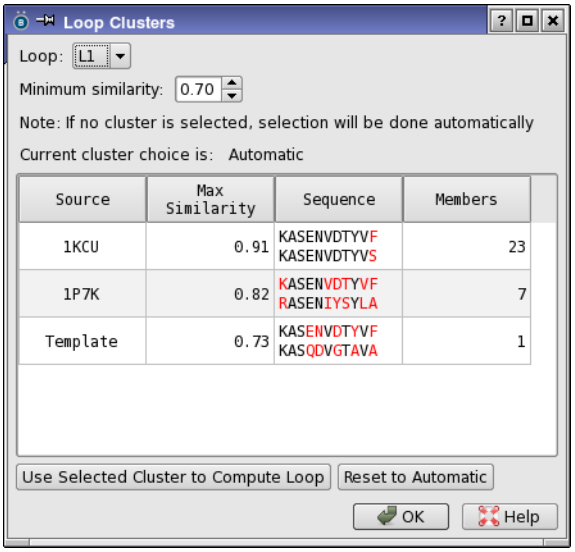

Figure 17.7. The Loop Clusters dialog box.

When this dialog box is first opened, the antibody databases are searched for loops of the same length as those in the query sequence for each of the six loops, and the loops are clustered structurally. The dialog box may therefore take a short while to open. Subsequently these loops and clusters are reused, so opening is much faster.

The clusters for the loop chosen from the Loop option menu are listed in the table, in order of decreasing similarity to the query sequence. The list is restricted to the clusters that have a loop whose similarity to the query loop is greater than the cutoff specified in the Minimum similarity box. The table columns are described in Table 17.2. You can sort the table by clicking on a column heading.

Table 17.2. Columns of the cluster table.

Column	Description
Source	PDB ID of the member of the given cluster that has the highest similarity to the query. The template used for the framework region is also included in the list.
Max Similarity	Similarity of the member of the cluster that has the highest similarity to the query.
Sequence	Sequence of the query for this loop (top) and of the most similar cluster member (bottom). Residues that differ are marked in red.
Members	Number of members of the cluster.

You can set the minimum acceptable similarity between a loop from the database and the query loop in the Minimum similarity box. Only the clusters that contain at least one loop whose similarity to the query is greater than this threshold are used in building the model, whether in the default automatic procedure or by manual selection of a cluster.

If the minimum cluster similarity is changed, this change applies to all loops. If no cluster is found with a loop that has the minimum similarity, then the program automatically uses the template loop in the database that has the greatest similarity to the query, no matter which cluster it belongs to.

To choose a loop cluster for a particular loop:

1. Select the loop from the Loop option menu.
2. Select the cluster in the table.
3. Click Use Selected Cluster to Compute Loop.

The cluster choice is shown above the table.

To clear a loop cluster selection:

1. Select the loop from the Loop option menu.
2. Click Reset to Automatic.

17.1.7 Refining Loops

Frequently, the model based on homology is adequate, and no further prediction is needed. However, in some cases, prediction of certain loops (especially H3) can be problematic, and advanced methods can result in better predictions.

The H3 loop, however, is often difficult to predict, and may need more advanced refinement. You can carry out advanced loop refinement using the Advanced Loop Model tab. From this tab, you can use Prime protein refinement to re-predict selected loops. If you generated more than one model, you should select the model you want to use from the Model to view option menu before leaving the Basic Loop Model tab.

Here are some general guidelines for when to refine a loop with Prime in the Advanced Loop Model tab:

- If the loop is long (9 residues or more) and the homology to the template is good (similarity above about 80%) then refinement with Prime is not usually necessary.
- If the loop sequence similarity is less than 40%, the basic model quality usually is very poor and a Prime refinement is recommended.

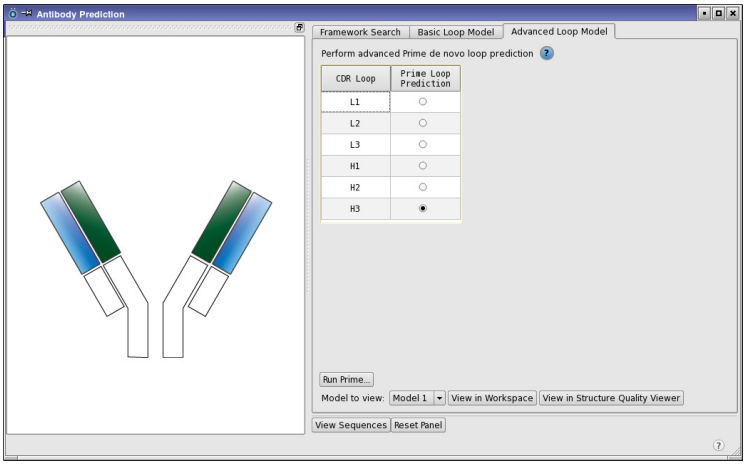

Figure 17.8. The Antibody Prediction panel, Advanced Loop Model tab.

- When building a new H3 loop with new sequences on the native (crystal) structure of an antibody, Prime refinement is recommended.

The quality of the Prime refinement is greater for shorter loops. Accurate and detailed loop predictions using Prime can take hours for each loop. It is not usually necessary to run an advanced loop prediction for any of the loops except the H3 loop. The H3 loop is selected by default for an advanced prediction, as it is the hardest loop to predict. The results are returned to the Project Table as new structures.

The input structure for the loop prediction is always the structure that came from the Basic Loop Model tab. This means that you cannot do sequential loop predictions in this panel, but you can use the Refinement panel to predict more than one loop (Tasks → Loop + Sidechain Prediction). See Chapter 6 of the *Prime User Manual* for more information on refinement tasks.

To choose the loop for prediction, select it in the table. Click Run Prime to run a Prime loop prediction job for the selected loop. A dialog box opens, in which you can name the job, select a host, and specify the maximum number of processors to use. The single, best loop prediction is returned. The Prime loop prediction algorithm is automatically selected based on the length of the loop.

The same controls as in the Basic Loop Model tab are present for viewing the structure.

17.1.8 Summary

The basic procedure for running a prediction using a homology model is as follows:

1. Import the sequences for the light and heavy chains: Click the part of the diagram for the region you want to import, and choose a source from the list that is displayed.
2. Set up the antibody databases that you want to use to search for homologs for the framework region and the loops:
3. (Optional) Click Choose Databases in the Framework Search tab to add databases and select them in the table. The default database is the one from the installation.
4. (Optional) Click Filter Search Results to filter the databases by properties of the structures in the databases.
5. Select Homology search and run the homology search by clicking Search.
6. Select the homolog in the results table that you want to use for the framework region, and click Accept.
7. (Optional) In the Basic Loop Model tab, select the loop cluster you want to use for each loop in the model by clicking Select Cluster, and choosing a cluster in the Loop Clusters panel. By default a cluster is chosen automatically, based on cluster size and a minimum similarity criterion.
8. (Optional) Include the antigen, ligands, cofactors, and water in the model.
9. (Optional) Set the number of models of the antibody that you want to generate.
10. Click Generate Loop Models to generate the models.
11. (Optional) If you think that the H3 loop needs further refinement, select it in the Advanced Loop Model tab, and click Run Prime.

The models are added to the Project Table.

You can also use coordinates from existing structures instead of homology models, for example if you want to vary just one of the loops. The procedure is similar to that given above.

1. Import the sequences for the light and heavy chains: Click the part of the diagram for the region you want to import, and choose a source from the list that is displayed.
2. Set up the antibody databases that you want to use to search for homologs for the framework region and the loops:
3. (Optional) Click Choose Databases in the Framework Search tab to add databases and select them in the table. The default database is the one from the installation.
4. Select Input coordinates and click Accept.

5. (Optional) In the Basic Loop Model tab, select the loop cluster you want to use for each loop in the model by clicking Select Cluster, and choosing a cluster in the Loop Clusters panel. By default a cluster is chosen automatically, based on similarity, then cluster size.
6. (Optional) Check the boxes for the loops for which you want to use input coordinates.
7. (Optional) Set the number of models of the antibody that you want to generate.
8. Click Generate Loop Models to generate the models.
9. (Optional) If you think that the H3 loop needs further refinement, select it in the Advanced Loop Model tab, and click Run Prime.

17.2 Humanizing Antibodies by Residue Mutation

If you have an antibody model that is not based on a human species, you can humanize it by mutating residues to create a more human-compatible protein. Appropriate sites for mutation can be found by comparing the antibody to homologs and identifying residues that satisfy various criteria, such as solvent accessibility and maximum number of side-chain interactions with the antigen or with the antibody itself. Multiple mutations can be done at each site, but only one site is mutated in each mutant structure.

The mutations can be set up and run in the Antibody Humanization: Residue Mutation panel, which you open by choosing Tasks → Antibody Modeling → Humanization → Residue Mutation in the main window. The panel has two tabs, one for setting up the criteria for choosing residues to mutate, and one for selecting the residues and defining the mutants. When both these tasks are done, you can start the job to mutate the residues.

17.2.1 Analyzing the Antibody

The first step in the process is to analyze the antibody to locate the antibody regions and calculate solvent-accessible surface areas (SASA). By default, the entire Workspace is analyzed, but if you want to limit the analysis, you can select Analyze only selected Workspace residues to limit the analysis to the Workspace selection. This is useful if you want to analyze the effect of the mutations on the binding of the antibody to an antigen, which can be done when the mutations are performed. You can include the entire complex in the Workspace and select only the antibody chains to analyze. If you want to examine binding, you must include the entire structure in the Workspace at this stage.

Once the analysis is complete, the sequence is shown in the panel's sequence viewer, colored by residue type, with its secondary structure assignment and disulfide bond annotation. You can choose which chain to display by using the Show in viewer options. This sequence is called the "parent" sequence.

17.2.2 Finding Homologs

The next step is to search for homologs in an antibody database. These homologs are used to identify residues that could be mutated, based on a comparison of the homologs with the parent sequence.

The path to the database that will be used for the search is displayed in the (noneditable) Antibody database text box. If you want to change the database, click Browse and navigate to the database. The default database is a database of human antibody data.

To find the homologs, click Search Antibody Database for Homologs. The progress of the search is shown in the status area at the bottom of the panel. When the search is done, the homologs must be aligned to the parent sequence, so that selection of residues for mutation can be done on the basis of matching residue positions. Click Align Homologs to perform the multiple sequence alignment of the homologs to the parent.

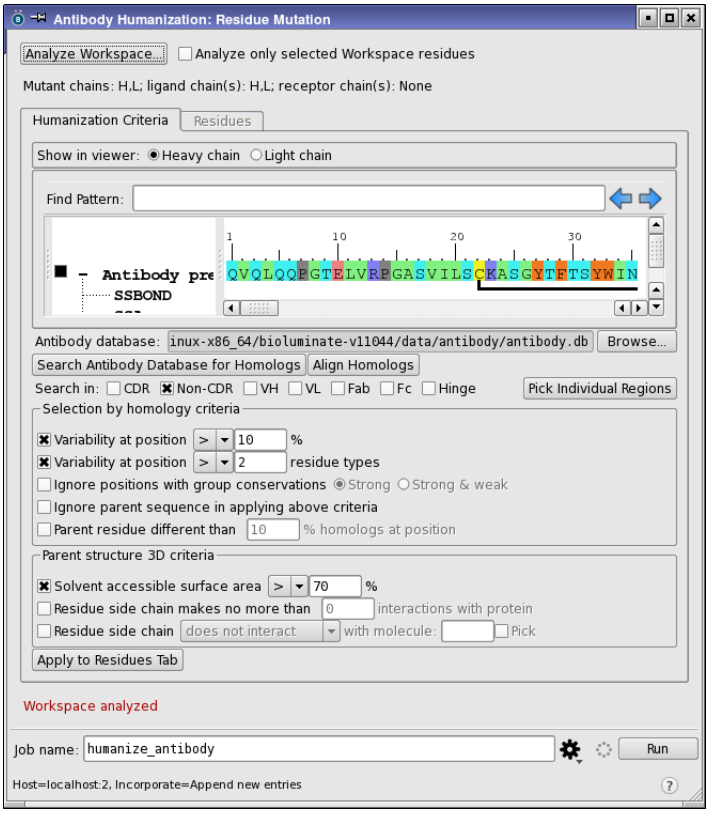

Figure 17.9. The Humanize Antibodies panel, Humanization Criteria tab.

17.2.3 Selecting Regions to Mutate

If you want to limit the mutations to specified regions of the antibody, you can do so by selecting one or more of the Substitute residues in options. The regions can be selected by group: CDR, Non-CDR, VH, VL, Fab, Fc, or Hinge. These options are available by default. For finer control, you can select individual regions. To display the options for the individual regions (which are hidden by default), click Pick Individual Regions. You can then pick the individual regions of the antibody for mutation. Note that you can only pick by group or pick individual regions: there is no connection between the two.

17.2.4 Setting Up Residue Selection Criteria

The next task is to set up criteria for selecting the residues to mutate. There are two kinds of criteria that you can use: those based on homology, and those based on the 3D structure of the parent.

Criteria that are based on homology use the variations in the residues at each residue position among the homologs, or between the homologs and the parent. The variations found are used as a basis for choosing default mutations. The criteria that you can set are:

- Variability at position >/< *N* %—Filter residues based on the percentage variability at the residue position. Choose whether to apply a minimum or a maximum variability from the option menu, and specify the percentage threshold in the text box.
- Variability at position >/< *N* residue types—Filter residues based on the residue type variability at the residue position. Choose whether to apply a minimum or a maximum variability from the option menu, and specify the threshold for the number of residue types that can vary in the text box.
- Ignore positions with group conservations—Ignore (do not select) residues that are strongly conserved or that are either strongly or weakly conserved. Select the appropriate option for strong conservation or both strong and weak conservation.
- Ignore parent sequence in applying above criteria—Ignore the parent (query) sequence when applying the variability or conservation criteria. Only the variability in the homologs is considered.
- Parent residue different than *N* % homologs at position—Select residues for which the parent residue is different from more than the specified percentage of homologs at the residue position.

The criteria based on the 3D structure of the parent antibody include solvent-accessible surface area (SASA) and interactions between residues. Interactions are determined by a distance

cutoff: any residue that has atoms within 4 Å of a given side chain is considered to interact with it.

- Solvent accessible surface area—Select this option to filter residues by their solvent-accessible surface area (SASA) relative to an isolated residue of the same type, and set a threshold for the maximum or minimum allowed relative SASA. This option is useful for locating surface (or buried) residues.
- Residue side chain makes no more than *N* interactions with protein—Select this option to filter out residues whose side chains make multiple interactions with the protein, and set the maximum number of such interactions.
- Residue side chain interacts/does not interact with molecule *N*—Select this option to filter residues by their interaction with a selected molecule. Choose whether to allow or disallow the interaction from the option menu, and specify the molecule in the text box, or select Pick and pick the molecule in the Workspace.

17.2.5 Selecting the Residues and Their Mutations

When you have finished setting up the selection criteria, click Select in the Residues Tab. The residues that match the criteria you provided in the Humanization Criteria tab are now selected in the Residues for consideration table in the Residues tab, and a set of mutations based on the variations seen in the homologs is chosen for these residues.

The Residues for consideration table lists all the residues in the Workspace. The first column identifies the residues. The second column specifies mutations of the residues. Any residue for which mutations are defined is mutated when the job is run, unless you select a subset of residues and select Mutate only selected residues.

To define the mutations for a given residue, click in the Mutations column. A menu is displayed in the column, from which you can select one or more residues to mutate to, or select groups of residue types. When you make a selection from the list, the residue is checked in the list and is added to the text box at the top of the menu. When you have finished selecting mutations, press ENTER. To cancel the current changes to the mutations, press ESC. To close the list, click the arrow button at the right of the text box.

To define a common set of mutations for multiple residues, first select the residues in the table. Then choose the residues from the For all selected rows, set mutations to menu, and click Apply. This menu works in the same way as the menu in the Mutations column of the table.

The number of mutations is reported at the bottom of the tab.

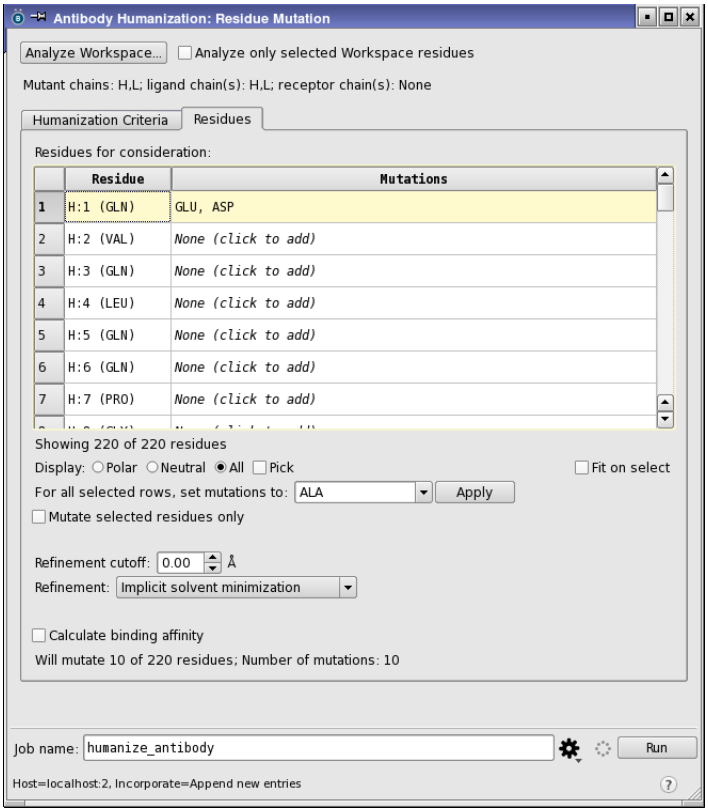

Figure 17.10. The Humanize Antibodies panel, Residues tab.

17.2.6 Setting Refinement Options for Mutated Residues

In the job that performs the mutations, the structure of the mutated residue is optimized, but you can also refine the structures of adjacent residues. The residues that are included in the refinement are selected by distance from the mutated residue, which is specified in the Cutoff text box. Any residue that has an atom within the specified distance of a hypothetical Arg residue at the mutation site is included in the refinement. A hypothetical Arg residue is used to ensure that the set of residues refined is identical regardless of the initial or mutated residue identities.

Before the minimization of the region around the mutation site, a side-chain prediction of the mutated residues is performed, which does a thorough exploration of possible side-chain conformations. The method used for the side-chain prediction can be chosen from the Refinement option menu. The choices are:

- Side-chain prediction—perform only the side-chain prediction.
- Side-chain prediction with backbone minimization—perform the side-chain prediction with minimization of the backbone atoms.
- Side-chain prediction with Cbeta sampling—perform the side-chain prediction including sampling of the CA-CB orientation.
- Side-chain prediction with backbone sampling—perform the side-chain prediction including sampling of the backbone conformations.

As part of the refinement, you can calculate the binding affinity of the antibody to the rest of the Workspace structure, by selecting Calculate binding affinity. The calculation is performed with the MM-GBSA method in the Prime program. This option is not available with gas phase minimization.

17.2.7 Running the Mutation Job

When you have finished making settings, you can run the job:

- To run the job with the current job settings, enter a name in the Job Name dialog box and click Run.
- To make job settings in the Job Settings dialog box, including the host, job name, and treatment of output, click the Settings button. Click Save and Run to run the job.

17.2.8 Analyzing the Results

When the job finishes, the Residue Scanning Viewer opens automatically, with the results loaded. You can also open this panel by choosing Tasks → Antibody Modeling → Humanization → Mutation Results, and then clicking Import in the panel to locate and load a set of results. The viewer is the same as for residue scanning, as the humanization job is just a residue scanning job set up for antibody mutations.

The Mutations table lists all the mutations that were generated, along with the changes in a range of properties as a result of the mutation. The properties include SASA (total, polar, and nonpolar), pK_a, hydropathy, number of rotatable bonds, energy, potential energy, and stability. These quantities are defined in Table 15.1 on page 109. You can sort the table columns by clicking on the column heading. You can plot any of these properties against the mutation (table row) by choosing the property from the Graph property option menu. If you want to export the table data as a CSV file, click Export, and navigate to a location and name the file.

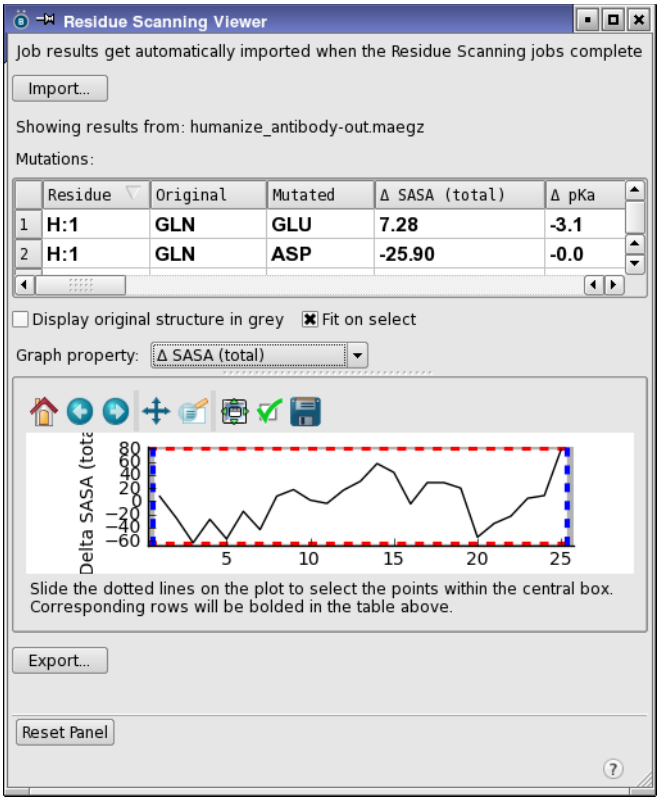

Figure 17.11. The Residue Scanning Viewer panel.

You can select a region in the graph using the horizontal and vertical dashed lines, which can be dragged to create the selection. The rows corresponding to the selected region of the graph are highlighted in the table above, and the residues are highlighted in the Workspace.

If you select a table row, the view zooms in to the mutated residue. To display the original structure, select Display original structure in grey. The parent antibody is displayed and colored grey. You can then see how the mutation is positioned in relation to the original residue.

17.2.9 Summary

1. Display the structure in the Workspace and analyze it, to identify residues and antibody features (Analyze Workspace).
2. Load an antibody database that is used to search for homologs.
3. Run the search (Search Antibody Database for Homologs) and align the homologs (Align Homologs).
4. (Optional) Select an option for the chain to show, and examine the alignment for that chain.
5. Choose the regions that you want to substitute residues in.
6. Specify the criteria for automatic selection of residues, in the Selection by homology criteria and Parent structure 3D criteria sections.
7. Click Apply to Residues Tab to select the residues that meet the criteria.
8. In the Residues tab, make any changes to the mutations.
9. Enter a name in the Job name text box and click Run, or click the Settings button, make settings in the Job Settings dialog box, and start the job.

   The mutated structures are incorporated into the project as new entries, and the Residue Scanning Viewer opens.

17.3 Humanizing Antibodies by CDR Grafting

Another way of creating humanized antibodies is to graft the CDR loops of a murine antibody onto a human framework. This task can be performed in the Antibody Humanization: CDR Grafting panel, which you open by choosing Tasks → Antibody Modeling → Humanization → CDR Grafting in the main window.

The human frameworks can be automatically located by a database search that finds the best matches to the query antibody. If you want to choose the databases to search for human antibodies, click Choose Databases. See Section 17.1.2 on page 121 for details of this task. All databases that you select as active are searched when you run a search for human frameworks.

The first task is to import the antibody structure. Click one of the Import antibody structure to humanize buttons to import the structure into the workflow:

- From Workspace—Use the structure that is displayed in the Workspace.
- PDB ID—Use the structure from the specified PDB ID, which you enter in the Enter PDB ID dialog box. The structure is retrieved from a local copy of the PDB or from the RCSB web site, depending on the preference set for PDB retrieval.
- Browse—Open a file browser in which you can navigate to the desired location and select the file that contains the antibody.

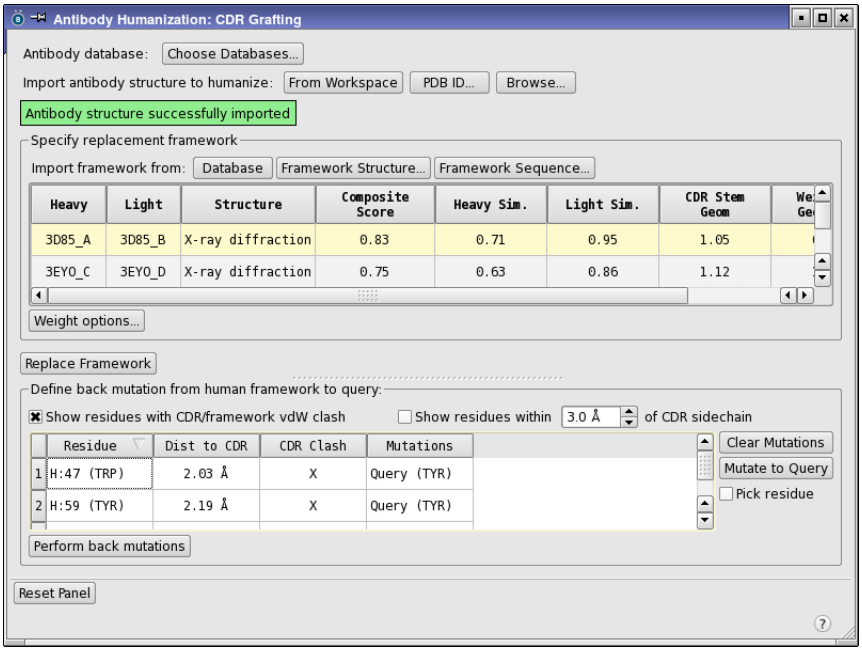

Figure 17.12. The Antibody Humanization: CDR Grafting panel.

Next, a structure for the framework is needed, which you import in the Specify replacement framework section. Click one of the Import framework from buttons to import a framework:

- Database—Search the antibody databases for the best-fitting human antibody frameworks. The human antibodies are aligned to the query antibody to find the best alignment of the framework regions to the query framework regions where the CDR loops join.
- Framework Structure—Import a file that contains the framework structure. The structure does not have to be an entire antibody, provided it contains the framework region.
- Framework Sequence—Import the sequence for the framework region. A homology model is built for the framework region using the default options, as represented in the Antibody Prediction panel.

The templates that were found or imported for the framework regions are listed in the Framework table, along with sequence identity and stem geometry scores. The results are sorted by score, which you can use to select a framework. You can also adjust the score by selecting the weights of its components in the Weights dialog box, which you open by clicking Weight Options.

To replace the original framework with the framework that is selected in the table, click Replace Framework. When the resulting model is build, a sequence viewer panel opens, showing the sequences of the original light and heavy variable chains and the chains of the grafted model, annotated with disulfide bridges, SSA, and CDR regions. This sequence viewer allows you to compare the original and the replacement framework.

The framework replacement is done without adjusting the loops in the CDR regions. There may be clashes between the CDR loops and the new framework, which can be relieved by back-mutation to the query, or to some other residue type.

The residues that clash can be listed in the table in the Define back mutation from human framework to query section, by selecting the options Show residues with CDR/framework vdW clash option or Show residues within *N* of CDR side chain (or both). You can adjust the threshold distance to the side chain if you wish.

Each framework residue that clashes is shown, along with the distance to the CDR side chain and a check for CDR clashes. The Mutations column lists the residue that the clashing residue will be mutated to.

- To select the residue you want to mutate to, click in the column and choose from the menu that is displayed. In addition to the standard residues, this menu has an item for the corresponding query residue and an item for no mutation, which disables mutation.
- To disable mutation of multiple residues, select them in the table and click Clear Mutations.

Table 17.3. Columns of the Framework table.

Column	Description
Heavy	PDB ID of the template used for the heavy framework region
Light	PDB ID of the template used for the light framework region
Structure	Source of the structure used for the framework regions, which can either be a crystal structure or a homology model. The scoring depends on the structure source.
Composite Score	Average of the Heavy Sim. and Light Sim. scores. This score is used as the sequence similarity score in the calculation of the weighted score.
Heavy Sim.	Sequence similarity of the entire variable domain sequence (framework and CDR) for the heavy chain. The similarity is the number of matching residues divided by the total number of residues, where "matching" means that the two residues have a positive score in the BLOSUM62 matrix.
Light Sim.	Sequence similarity of the entire variable domain sequence (framework and CDR) for the light chain. The similarity is the number of matching residues divided by the total number of residues, where "matching" means that the two residues have a positive score in the BLOSUM62 matrix.
CDR Stem Geom	Fitness score for the geometry of the CDR loop stem residues. The stem residues are the residues in the framework region that are directly attached to the CDR loops. The geometry score is related to the RMSD between the native and the grafted antibody for the C-N distance, the C-N-C angle, and the C-C-N-C dihedral across the bond between the stem residue and the first (or last) loop residue.
Weighted Geom+Sim	Weighted combination of the stem geometry fitness score and the sequence similarity score. The rows in the table are ordered by this score. The weights can be specified by clicking Weight Options and specifying the weights in the Weights dialog box.

- To set the mutation to the query for multiple residues, select them in the table and click Mutate to Query.
- To select residues in the table by picking them in the Workspace, select Pick residue, then pick the residues. This allows you to use the 3D structure to find residues to mutate.

When you have decided which residues to mutate and what to mutate them to, click Perform Back Mutations. The residues shown in the table that have a valid mutation in the Mutations column are mutated, and the mutated residues and neighbors within 4 Å are refined.

After you have performed these tasks, you may want to refine the structure of the antibody further to relieve strain. For example, you might want to do a side-chain refinement on the framework to relieve minor clashes with the CDR loops.

17.4 Antibody Databases

Modeling antibodies is done with the help of a database of prepared antibody structures. A database is supplied with the distribution, and is the default database. You can create and manage your own databases in the Antibody Database Management, which you open by choosing Tasks → Antibody Modeling → Database Management in the main window.

Structures that are added to a database are automatically characterized and curated. Antibody structures and the light and heavy chains are identified using a similarity search against known antibodies, and structures that do not meet the similarity criteria are rejected. Then the constituent regions of the antibody chains, including the framework region, as well as the six hypervariable loops, are identified and annotated for use in subsequent predictions.

When you first open the panel, the default database is loaded. Normally this database is installed by an administrator and you do not have write privileges, so it is opened read-only. You can import this database into your own database if you want, as described below. It is only necessary to do this if you want to modify the data in some way, because you can specify multiple databases for modeling.

To open a database, or to create a new database, click Open/Create. A file selector opens, and you can navigate to the desired location. If you want to open a database, select the database from the list of files. It should have the extension `.db`. If you want to create a new database, enter the name in the File name text box.

The structures in the database are shown in the antibody table. You can restrict the structures that are listed in the table by searching for strings in the table and only showing the rows that contain the string. By default, all visible text is searched, but you can change it by clicking

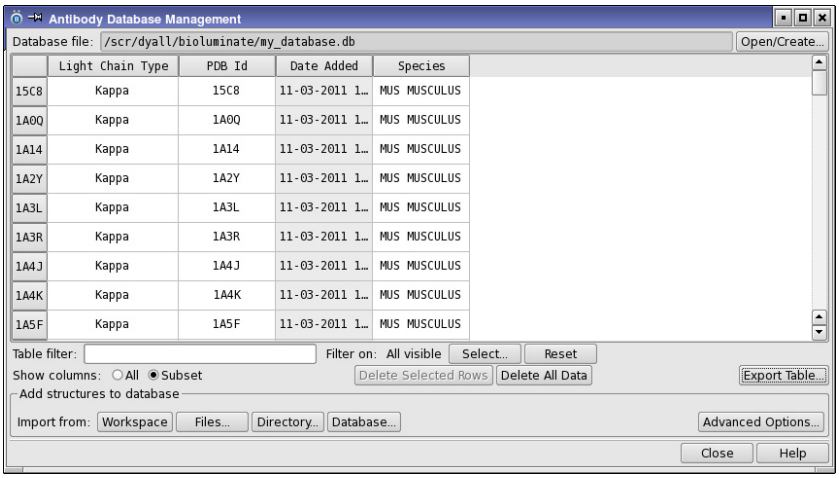

Figure 17.13. The Antibody Database Management panel.

Select and choosing the columns you want to search on. To return to sorting all visible columns, click Reset. The text box has a tool tip that explains the syntax for the search string, which can include relational operators, wild cards, regular expressions, and some special terms.

The table shows only a few columns by default. If you want to see all the columns, by Show columns, select All. The full set of columns includes the identity (residue range) and length of each loop, and a range of information from the originating PDB structure. If you want to export the information in the columns to a CSV file, click Export Table, and navigate to a location and name the file in the file selector that opens.

To delete structures from the database, select them in the table and click Delete Selected Rows. If you want to clear the entire database, click Delete All Data. You should exercise care when deleting rows or all data, as these functions are not reversible.

Structures can be added to the database from several sources, represented by buttons in the Add structures to database section. In each case, the imported data is automatically processed for you: the antibody chains are identified, the constituent regions of the chains are determined, and all the pertinent information required for subsequent modeling is saved in a rapidly accessed format. You need only supply the antibody structures in PDB (or Maestro) format.

- Workspace—import structures from the Workspace.
- Files—import structures from PDB files or Maestro files. Opens a file selector, in which you can navigate to and select multiple files for import.
- Directory—import all structures from a directory. Opens a directory selector, in which you can navigate to the directory to import from.
- Database—import structures from an existing database. Opens a file selector, in which you can navigate to and select the database (.db) to import from. Import progress is displayed in a bar at the bottom of the panel.

The structures are filtered to ensure that only those structures that have characteristics of antibodies are included in the database. You can alter the criteria for filtering structures in the Antibody Database - Advanced Options dialog box, which you open by clicking Advanced Options. You can also select the file types that are presented when importing structures from files.

In the Minimum sequence identity section, you can specify cutoffs for determining whether a structure should be included in the database, based on percentage sequence identity. If no chain in the structure passes the tests, it is not included in the database. A chain must meet the FR threshold and either the VL or VH threshold to be included in the database. If a chain meets only one of the VL or VH thresholds, it is only included as a light or heavy variable chain.

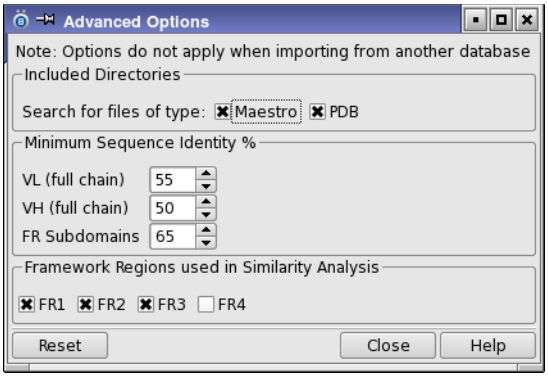

Figure 17.14. The Antibody Database Management Advanced Options panel.

For the framework regions, you can specify which of the four framework regions to use in the similarity analysis, by selecting the options for FR1, FR2, FR3, or FR4 in the Framework regions used in similarity analysis section.

To revert to the default options, click Reset.

References

1. Olsson, M. H. M.; Søndergard, C. R.; Rostkowski, M.; Jensen J. H. PROPKA3: Consistent Treatment of Internal and Surface Residues in Empirical pK_a predictions. *J. Chem. Theor. Comput.*, **2011**, *7*, 525–537.

2. North, B.; Lehmann, A.; Dunbrack, R. L., Jr. A new clustering of antibody CDR loop conformations. *J. Mol. Biol.* **2011**, *406*, 228.

3. Comeau, S. R.; Gatchell, D. W.; Vajda, S.; Camacho, C. J. ClusPro: an automated docking and discrimination method for the prediction of protein complexes. *Bioinformatics* **2004**, *20*, 45–50.

4. Comeau, S. R.; Gatchell, D. W.; Vajda, S.; Camacho, C. J. ClusPro: a fully automated algorithm for protein–protein docking. *Nucleic Acids Res.* **2004**, *32*, W96–W99.

5. Kozakov, D.; Brenke, R.; Comeau, S. R.; Vajda, S. PIPER: an FFT-based protein docking program with pairwise potentials. *Proteins* **2006**, *65*, 392–406.

6. Kyte, J.; Doolittle R.F. A simple method for displaying the hydropathic character of a protein. *J. Mol. Biol.* **1982**, *157*, 105–132.

7. Hellberg, S.; Sjöström, M.; Skagerberg, B.; Wold, S. Peptide quantitative structure-activity relationships, a multivariate approach. *J. Med. Chem.* **1987**, *30*, 1126.

8. Sandberg, M.; Eriksson, L.; Jonsson J.; Sjöström, M.; Wold, S. New chemical descriptors relevant for the design of biologically active peptides. A multivariate characterization of 87 amino acids. *J. Med. Chem.* **1998**, *41*, 2481.

9. Tian, F.; Lv, F.; Zhou, P.; Yang, Q.; Jalbout, A. F. Toward prediction of binding affinities between the MHC protein and its peptide ligands using quantitative structure-affinity relationship approach. *Protein Pept. Lett.* **2008**, *15*, 1033.

10. Lawrence, M. C.; Colman, P. M. *J. Mol. Biol.*, **1993**, *234*, 946–950.

Getting Help

Information about Schrödinger software is available in two main places:

- The `docs` folder (directory) of your software installation, which contains HTML and PDF documentation. Index pages are available in this folder.
- The Schrödinger web site, http://www.schrodinger.com/, particularly the Support Center, http://www.schrodinger.com/supportcenter, and the Knowledge Base, http://www.schrodinger.com/kb.

Finding Information in Maestro

Maestro provides access to nearly all the information available on Schrödinger software.

To get information:

- Pause the pointer over a GUI feature (button, menu item, menu, ...). In the main window, information is displayed in the Auto-Help text box, which is located at the foot of the main window, or in a tooltip. In other panels, information is displayed in a tooltip.

  If the tooltip does not appear within a second, check that Show tooltips is selected under General → Appearance in the Preferences panel, which you can open with CTRL+, (⌘,). Not all features have tooltips.

- Click the Help button in the lower right corner of a panel or press F1, for information about a panel or the tab that is displayed in a panel. The help topic is displayed in your browser. The button may have text or an icon:

  

- Choose Help → Online Help or press CTRL+H (⌘H) to open the default help topic in your browser.
- When help is displayed in your browser, use the navigation links or search the help in the side bar.
- Choose Help → Manuals Index, to open a PDF file that has links to all the PDF documents. Click a link to open the document.

- Choose Help → Search Manuals to search the manuals. The search tab in Adobe Reader opens, and you can search across all the PDF documents. You must have Adobe Reader installed to use this feature.

For information on:

- Problems and solutions: choose Help → Knowledge Base or Help → Advanced Help Options → Known Issues → *product*.
- Software updates: choose Help → Check for Updates.
- New software features: choose Help → Advanced Help Options → New Features.
- Scripts available for download: choose Tasks → Scripts → Update.
- Python scripting: choose Help → Advanced Help Options → Python Module Overview.
- Utility programs: choose Help → About → About Utilities.
- Keyboard shortcuts: choose Help → Keyboard Shortcuts.
- Installation and licensing: see the *Installation Guide*.
- Running and managing jobs: see the *Job Control Guide*.
- Using Maestro: see the *Maestro User Manual*.
- Maestro commands: see the *Maestro Command Reference Manual*.

Contacting Technical Support

If you have questions that are not answered from any of the above sources, contact Schrödinger using the information below.

E-mail: help@schrodinger.com
USPS: Schrödinger, 101 SW Main Street, Suite 1300, Portland, OR 97204
Phone: (503) 299-1150
Fax: (503) 299-4532
WWW: http://www.schrodinger.com
FTP: ftp://ftp.schrodinger.com

Generally, e-mail correspondence is best because you can send machine output, if necessary. When sending e-mail messages, please include the following information:

- All relevant user input and machine output
- BioLuminate purchaser (company, research institution, or individual)
- Primary BioLuminate user
- Installation, licensing, and machine information as described below.

Gathering Information for Technical Support

The instructions below describe how to gather the required machine, licensing, and installation information, and any other job-related or failure-related information, to send to technical support. Where the instructions depend on the profile used for Maestro, the profile is indicated.

For general enquiries or problems:

1. Open the Diagnostics panel.
   - **Maestro:** Help → Diagnostics
   - **Windows:** Start → All Programs → Schrodinger-2013-3 → Diagnostics
   - **Mac:** Applications → Schrodinger2013-3 → Diagnostics
   - **Command line:** `$SCHRODINGER/diagnostics`
2. When the diagnostics have run, click Technical Support.

   A dialog box opens, with instructions. You can highlight and copy the name of the file.
3. Attach the file specified in the dialog box to your e-mail message.

If your job failed:

1. Open the Monitor panel.
   - **Maestro:** Applications → Monitor Jobs or Tasks → Monitor Jobs
   - **BioLuminate:** Tasks → Job Monitor
   - **Elements:** Tasks → Monitor Jobs
2. Select the failed job in the table, and click Postmortem.

   The Postmortem panel opens.
3. If your data is not sensitive and you can send it, select Include structures and deselect Automatically obfuscate path names.
4. Click Create.

   An archive file is created in your working directory, and an information dialog box with the name of the file opens. You can highlight and copy the name of the file.
5. Attach the file specified in the dialog box to your e-mail message.
6. Copy and paste any log messages from the window used to start the interface or the job into the email message, or attach them as a file.
   - **Windows:** Right-click in the window and choose Select All, then press ENTER to copy the text.

- **Mac:** Start the Console application (Applications → Utilities), filter on the application that you used to start the job (Maestro, BioLuminate, Elements), copy the text.

If Maestro failed:

1. Open the Diagnostics panel.
   - **Windows:** Start → All Programs → Schrodinger-2013-3 → Diagnostics
   - **Mac:** Applications → SchrodingerSuite2013-3 → Diagnostics
   - **Linux/command line:** `$SCHRODINGER/diagnostics`
2. When the diagnostics have run, click Technical Support.

   A dialog box opens, with instructions. You can highlight and copy the name of the file.
3. Attach the file specified in the dialog box to your e-mail message.
4. Attach the error files your e-mail message.

   The files should be in the following location:

- **Windows:** `%LOCALAPPDATA%\Schrodinger\appcrash`
  (Choose Start → Run and paste this location into the Open text box.)
  Attach `maestro_error_`*`pid`*`.txt` and `maestro.exe_`*`pid_timestamp`*`.dmp`.
- **Mac:** `$HOME/Library/Logs/CrashReporter`
  (Go → Home → Library → Logs → CrashReporter)
  Attach `maestro_error_`*`pid`*`.txt` and `maestro_`*`timestamp_machinename`*`.crash`.
- **Linux:** `$HOME/.schrodinger/appcrash`
  Attach `maestro_error_`*`pid`*`.txt` and `crash_report_`*`pid`*`.txt`.

120 West 45th Street
17th Floor
New York, NY 10036

101 SW Main Street
Suite 1300
Portland, OR 97204

245 First Street
Riverview II, 18th Floor
Cambridge, MA 02142

8910 University Center Lane
Suite 270
San Diego, CA 92122

155 Gibbs St
Suite 430
Rockville, MD 20850-0353

Dynamostraße 13
D-68165 Mannheim
Germany

Zeppelinstraße 73
D-81669 München
Germany

Potsdamer Platz 11
D-10785 Berlin
Germany

Quatro House
Frimley Road
Camberley GU16 7ER
United Kingdom

8F Pacific Century Place
1-11-1 Marunouchi
Chiyoda-ku, Tokyo 100-6208
Japan

No. 102, 4th Block
3rd Main Road, 3rd Stage
Sharada Colony
Basaveshwaranagar
Bangalore 560079, India

SCHRÖDINGER®